2015 XSEDE Federation Risk Assessment Overview

01: Risk Assessment Approach & Broader Goals

This document has been assembled from documentation on the XSEDE Wiki, originally created in 2012 and later revised in 2015. It captures the work associated with performing the 2012 and 2015 XSEDE Risk Assessments and forms the basis of future XSEDE risk assessments.

The risk assessment methodology used is the NIST 800-30 standard. The justification for this approach and how it fits in the larger risk management strategy is described in 02: A Risk-based Security Program for XSEDE.

The scope of this effort and the actors involved are further described in the 03: Risk Assessment Project Sizing.

Risk Assessment Team

The initial effort utilized a team with representation from each of the operations divisions: Security, Data Services, XSEDENet, Software Support, Accounting and Account Management, and System Operational Support. We needed a broad team to gather information, vet documents and help determine the impact of threats against key assets.

The update to this original assessment can be performed much more easily and only requires regular engagement from the XSEDE Security Working Group. Domain experts will be consulted as needed to determine the impact of key risks.

System Characterization

The system characterization defines the boundaries of investigation. It determines what hardware, software, services, interfaces and data are within that scope. It also discusses who uses the system, who supports it, and how it relates to XSEDE's mission.

See 04: XSEDE System Characterization for details.

Threat Profile

A threat is the potential for a particular threat source to exercise a particular vulnerability towards a malicious end. If there is no vulnerability, there is no threat. And a threat is not the person or event that triggers a threat; that is the threat source (e.g., hacker, inept programmer, or flood).

A threat profile is usually done before the vulnerability assessment. Even though you may not know specifically if you have vulnerabilities for the threats to exploit, you can state what are the expected common threat actors as well as what you are not going to protect against or address in the scope of this assessment. For us, we are ignoring natural and environmental threats, which would have to be determined separately one level lower by the major service providers. Instead, we focus on human threats, both intentional (e.g., direct attack or theft) and unintentional (e.g., poorly written software). We only consider threats to shared XSEDE infrastructure, XSEDE's trust fabric and the overall stability of the federation. Site specific threats should be considered in more detail by individual SPs.

See 05: Threat Profile for details.

Vulnerability Identification

A vulnerability is a flaw or weakness in system security procedures, design, implementation or internal controls that could potentially be exercised by a threat agent to result in a breach or violation of the system's security policy. In this phase of the risk assessment we identify as many federation-level vulnerabilities in XSEDE as we can, mapping each vulnerability to a threat action and threat source.

06: Vulnerability Identification organizes all of this into one large table. Information was originally gathered by investigating assets from 04: XSEDE System Characterization, reviewing other risk assessments, interviews with people from each operations domain, and review and acceptance by the risk assessment team.

The 2015 updates to this section were made based on changes to 04: XSEDE System Characterization, adding and retiring risks based on changes in XSEDE software and services. Further changes were made as changed security activities had mitigated certain vulnerabilities since the original assessment.

Control Analysis

Security controls are mechanisms in place to mitigate the risk of threats being realized and hence exploiting vulnerabilities in your infrastructure. Controls can be administrative (e.g., policies, standards, guidelines, training and other processes), technical/logical (e.g., authentication and authorization systems, file permissions, firewalls, intrusion detection systems, etc.), or physical (e.g., locked file cabinets, secured data centers, cameras, fences, etc.).

Since risks are addressed by controls, it is important to understand the security controls already in place or planned. Without this step, you don't know where there are gaps. One of the outputs of the risk assessment is to recommend changes and additions to controls to address the highest impact and most probable risks.

07: Control Analysis Matrix lists all the existing and planned controls with current status and notes. The final report will recommend modifications to this set of controls.

Likelihood Determination

In this step we qualitatively determine the likelihood that a vulnerability will be exploited within the scope of 05: Threat Profile derived earlier. We then rank the likelihood as high, medium or low. In this ranking three things are considered:

- the motivation and capabilities of the threat source
- the specifics of the vulnerability
- the effectiveness of current controls to mitigate associated risks

Borrowing a table from the NIST 800-30, here are the definitions for high, medium and low.

LIKELIHOOD LEVEL	LIKELIHOOD DEFINITION
High	The threat-source is highly motivated and sufficiently capable, and controls to prevent the vulnerability from being exercised are ineffective.
Medium	The threat-source is motivated and capable, but controls are in place that may impede successful exercise of the vulnerability.
Low	The threat-source lacks motivation or capability, or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised.

Each vulnerability, and hence risk, maps to a row in the table from 06: Vulnerability Identification. Rather than replicate the table in several places where it can become out of sync, we added new likelihood column to the existing table.

Decision Process

As a first round, we used blinded (to all but the risk assessment lead) votes via email for each item. As long as 2/3 of the team voted on an item and the votes had a standard deviation less than 0.7, we took the average of the votes and chose the closest probability ranking. To compute a numerical average, each ranking received a numerical value of 1, 2 or 3, and the average was simply the arithmetic mean.

For other items, there is either wide-spread disagreement or a lack of understanding to gain quorum. These were resolved over a phone meeting. If quorum was the problem and the team lead could not get enough votes, then a threat scenario would need to have been written up on the issue by the most expert person of the team for that item. Other items were discussed on a phone meeting with quorum until a majority decision was made (All items had quorum in this case). The call started with a reminder of what the different rankings of high, medium and low mean, and the chair's synopsis of the disputed items and why he voted how he did.

- Initial Risk Rankings Votes Matrix
- Final votes with phone meeting notes

2015 Updates

Based on changes in 04: XSEDE System Characterization and the set of security control changes since the original assessment, risks were updated by the security analyst with review and advisement from the larger XSEDE Security Working Group. Some risks were mitigated or retired based on changes in XSEDE, three new risks were added.

Impact Analysis

In the previous step we estimated the likelihood of particular vulnerabilities being exploited by specific threats. In this step we evaluate the impact of such exploitations. This is somewhat more challenging and fuzzy as we have to keep several things in mind.

First, the impact of any exploit is going to depend upon (1) the mission of XSEDE, (2) the criticality of the vulnerable system or data to XSEDE, and (3) the sensitivity of the affected system or data. This information is usually found in a mission or business impact analysis, but XSEDE has no such documentation. There is a mission statement noted in 04: XSEDE System Characterization, and there is

a categorization of XSEDE services into tiers in the XSEDE services master spreadsheet, however. But in cases where it is not clear how critical a system is, we relied upon the system owner or maintainer's help.

Impact from a security incident could affect the integrity, availability or confidentiality of a system or data. Depending on the system affected, we could be concerned more with one kind of impact than another. Particularly, in light of XSEDE's mission, it would most often be that integrity and availability are of more concern than confidentiality. However, the impact on each of these three properties should be considered for any potential exploit.

Borrowing a table from the NIST 800-30, here are the definitions for high, medium and low.

IMPACT LEVEL	IMPACT DETERMINATION
High	Exercise of the vulnerability (1) may result in the highly costly loss of major tangible assets or resources; (2) may significantly violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human death or serious injury.
Medium	Exercise of the vulnerability (1) may result in the costly loss of tangible assets or resources; (2) may violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human injury.
Low	Exercise of the vulnerability (1) may result in the loss of some tangible assets or resources or (2) may noticeably affect an organization's mission, reputation, or interest.

Decision Process

We follow the same process of voting on impact as in the previous phase of likelihood determination. The results of these votes are recorded in the same table on 06: Vulnerability Identification as the likelihood decisions are. The detailed spreadsheet results are linked to below.

- Final Risk Assessment Votes

2015 Updates

Based on changes in 04: XSEDE System Characterization and the set of security control changes since the original assessment, risks were updated by the security analyst with review and advisement from the larger XSEDE Security Working Group. Some risks were mitigated or retired based on changes in XSEDE, three new risks were added.

Risk Determination

With likelihood and impact analysis complete, we can calculate numeric values for each risk to rank them. We use the table below which gives a weight for both likelihood and impact. Numerical weights are in parentheses, and underneath are the number of risks of the give ranking in bold.

	LOW IMPACT (10)	MEDIUM IMPACT (50)	HIGH IMPACT (100)
High Likelihood (1.0)	Low (10) 0	Medium (50) 3	High (100) 0
Medium Likelihood (0.5)	Low (5) 6	Medium (25) 12	Medium (50) 0
Low Likelihood (0.1)	Low (1) 9	Low (5) 4	Low (10) 0

There are a few things to notice from this ranking. First, as we had no high impact risks, none resulted in a total risk labeled as high. The highest we had were 3 Medium risks with highly likely exploitation and medium impact. These 3 risks with a ranking of 50 are a natural place to start when considering what risks to address with new controls first.

Second, we see that the most common ranking is 25, a medium risk, but lower than the other three mediums. There are more low rated risks than all the others combined when we add together the three types of low risks. Work after addressing the three rank 50 risks should probably take from the pool of risks ranked 25. Resources and available controls will influence the priority among those 12 risks.

The four rank 50 risks are:

- that passwords can be easily compromised and harvested
- that a zero-day root escalation exploit is discovered in a common piece of software used by XSEDE and exploited by an attacker with stolen credentials
- that because there is no regular security training for users or service providers, policies are not followed and lead to mistakes or accidents that attackers can exploit

Specific rankings for each risk can be found in 06: Vulnerability Identification.

Control Recommendations

As noted above in the control analysis section above, XSEDE already employs many security controls to meet its operational and security goals. In this section we list recommended changes or additional controls to mitigate, contain, eliminate or transfer the risks discovered in this process. This are prioritized high, medium or low. High priority controls address the risks ranked 50 or above; medium priority addresses risks rated 25-49; and low is for risks rated less than 25.

In 08: Control Recommendations we organize all of these controls and note which risks they address by numerical index in 06: Vulnerability Identification.

Accepted Risks

Some of these controls are not likely worth the cost, and it is not merely an issue of priority. Therefore, not all of the controls are labeled as RECOMMENDED in 08: Control Recommendations. Risks with no recommended controls are one that we recommend XSEDE ACCEPT. In these cases, the cost of mitigating controls from usability, performance or raw dollars cannot be justified.

These are risks: #1, #4, #6, #8, #10, #15, #19, #20, #21, and #24

Download the 2015 XSEDE Risk Assessment Final Report.

02: A Risk-based Security Program for XSEDE

The Goal

Security Management is at the core of an organization's information security structure. As Shon Harris says, "Security management includes risk management, information security policies, procedures, standards, guidelines, baselines, information classification, security organization, and security education."[1] This is a continuous process, and the goal is to develop and help implement a formal, risk-based security program for XSEDE.

This process is guided by a security analyst, working with the XSEDE Security Officer and the XSEDE Security Working Group. A security analyst works at a higher, more strategic level than security admins and incident responders, and he helps develop policies, standards, and guidelines, as well as set various baselines. Whereas security admins, system administrators and incident responders are focused on daily operational needs and specialize on pieces and parts of the security program, a security analyst helps define the security program elements and follows through to ensure the elements are coordinated and carried out properly as a whole.

Planned Deliverables for XSEDE

To develop a formal risk-based security program for XSEDE several things were needed. All of these have been developed to some extent during the first 3 years of XSEDE, but they require periodic updating.

- a risk assessment for XSEDE at a federation level and a process for individual sites to drill down and perform similar assessments for their programs;
- federation security policies, procedures and security blueprints (i.e., logical/technical controls to mitigate federation-wide risks for XSEDE);
- a more well-defined security organizational structure with clear roles; and
- plans for educating users and raising security awareness

The Process

The first 2 major goals, (a risk assessment and a set of policies, procedures & security blueprints), are tightly coupled and were completed in a sequential order during the first year. They form the backbone of XSEDE's risk management strategy, which is the process of identifying risk, reducing risk and implementing controls. Our risk management process can be described at a high-level by the following steps, which feed into each other: Threat Profile to Risk Assessment to Security Plans to Policies, Procedures and Blueprints.

Threat Profile

This involves understanding who and what could damage federation assets, identifying types of attacks and crimes that could take place, and understanding the impact of these threats and threat actors. We focus on cyber security and only federation-level threats. Threats that are specific to certain sites or don't affect the federation as a whole are considered when individual sites carry out their own risk assessments. We consider the types of adversaries or threat actors at a very high-level here with their general capabilities. This threat profile tries to capture the overall level of acceptable risk to XSEDE, and

[1] Much of the language and definitions borrow from Shon Harris's CISSP study book.

hence required input from the leaders of XSEDE. While the XSEDE Security Working Group can inform this, the entire management team are stakeholders here.

There are many kinds of threats, including environmental, application error, physical failure and disasters. Our goal is more focused to prioritize cybersecurity resources and develop security blueprints. Therefore, we have focused on application errors and a subset of threats called inside/outside threats, i.e. hacking, cracking and attacks, when we created our threat profile and performed the follow-on risk assessment.

Information Risk Assessment

A crucial part of the risk management process is risk assessment. Before a risk assessment is done, you must understand the risk tolerance of the organization and the general characterization of threats. This is why a short (1-2 page) threat profile was completed first. It was also critical, to scope the problem and do a Risk Assessment Project Sizing before starting the project. The sizing was used to decide kinds of assets and threats considered, the depth of analysis, and what can be done with the resources we have to carry out the analysis. A team that comprises multiple stakeholders was formed for the initial risk assessment, though not for periodic updates.

While there are many types of risk analysis processes (e.g., NIST 800-30, OCTAVE, and AS/NZS 4360:2004), all processes can be broken down at a high-level to the following steps:

- identify assets & values;
- identify threats and potential vulnerabilities;
- determine impact/loss from realized threats and their likelihood of occurrence; and
- balance impact with countermeasure selection.

Risk assessments can be quantitative, qualitative, or a mix. Some things simply cannot be quantified, such as, reputation. Therefore, any complete risk assessment is likely to have at least some qualitative parts. It is particularly dangerous, and far too common, to make a quantitative assessment with fuzzy inputs. This leads to an impression of an exactness that is rarely there. Such a false sense of certainty often leads to bad decisions being made. Because of the difficulty getting quantitative measures for XSEDE (impact is often not financial but to reputation, loss of cycles, availability, etc) and the fact that the budget is fixed for security, we started with a qualitative approach. This especially makes sense for XSEDE, where the main goal of this risk assessment is to inform prioritization of limited resources to deploy the countermeasures that give the best cost benefit.

Significant threats were written up with a one page scenario by an expert (w.r.t. the asset being threatened) with guidance from the security analyst. These scenarios discussed how the threat could be carried out and the potential impact. Help was needed from all the stakeholders, especially to understand the impact of threats to assets and business processes. Predicting the likelihood of a threat being realized was easier for the analyst to do alone, since it is tied closely to an understanding of the vulnerabilities, but it was still reviewed by other technical domain leaders with experience operating the infrastructure as a reality check.

The NIST 800-30 process, CMU created OCTAVE and Australian AS/NZS 4360:2004 are common risk assessment methodologies. We used the NIST 800-30 process this for many reasons. First, each of these build upon the previous and no time is wasted by starting with the NIST method, expanding to the OCTAVE approach, and finally moving onto the Australian/New Zealand standard (now superseded by the ISO 31000 series). In fact, many professionals recommend that an organization get their feet wet with the NIST approach first simply because it is a lot to tackle everything at once to start with the deepest and most sophisticated approach. Second, the OCTAVE approach relies heavily upon a series of self-directed workshops with management, operations, security and business heads walking through several scenarios, questionnaires and checklists. This is a significant commitment from the whole organization that XSEDE was unlikely to make as a whole. And even if it did want to, it would have been hard to do with teams so geographically distributed. Therefore, we started with the NIST 800-30 process, which maps nicely on top of the high-level process described above.

Security Plans

Once the risk assessment was finished and we figured out which risks to mitigate (through specific countermeasures), accept, transfer or avoid, then we made security plans[2] for subsequent years based on the assessment.

One of the first actions was creating new policies for XSEDE. While policies, standards, baselines, guidelines and procedures are all a part of a security policy architecture, they are distinct. A security policy is an overall statement produced, or at least endorsed, by senior management which dictates the role security plays in the organization. These are very high-level and focus on high-level goals, roles and responsibilities and the acceptable level of risk in an organization. Policies should be very stable over time. Standards are mandatory actions, activities or rules. Standards describe how a policy is realized in practice and may refer to specific technologies or processes to be followed. Guidelines are similar to standards but are recommended. Finally, procedures are low-level how-to's for implementing the standards or guidelines.

The technical countermeasures selected in the final step of the risk assessment fed into the security blueprints. These blueprints lay out the solutions and components needed to fulfill XSEDE's security needs, and are reflected in several SDI activities, though without formally introducing the term security blueprint. In XSEDE, these blueprints did not come wholly from the risk assessment process but were also be influenced by what was promised to program officers and stakeholders. In general, it should be possible map to the controls in the blueprints to security controls in the ISO 17799 Part 1 or the newer ISO 27000 series, though a specific effort was not made to do this formally.

XSEDE does not have a formal security architecture, which is often developed for products seeking certain certifications like the NIST FIPS program. Such an architecture would connect the security policies formally to a set of mechanisms to realize the policy in a formal method.

[2] Mitigation is usually through technical countermeasures, and avoidance is usually through policy and procedure. Transferring risk, usually involves insuring against a threat, something XSEDE is very unlikely to do. Acceptance does not mean ignoring. It means a risk is acknowledged determined by senior management not viable to reduce or avoid.

03: Risk Assessment Project Sizing

Background & Purpose

We perform a threat and risk assessment for XSEDE to:

1. inform future security plans;
2. prioritize resources of the security operations team and its activities; and
3. demonstrate a process for service providers to use in their security planning.

Scope

A risk assessment can be scoped both in terms of threats and assets considered. We, the XSEDE security team, perform risk assessments at the federation-level and then leave it to individual Service Providers (SPs) to follow a similar process for their local XSEDE resources. We consider threats that affect more than one site, the fundamental shared resources of XSEDE, or the underlying trust fabric. For example, threats specific to the architecture of a particular HPC would not be considered here, but threats against an XSEDE authentication system (e.g., Kerberos KDC) would be.

The most general kind of risk assessment would consider all sorts of threats, including environmental, application error, physical failure and disasters. However, our goal is namely to wisely choose cyber security controls to give us the most bang for our buck. Therefore, we only consider cyber security threats, such as, application errors and insider/outsider threats (e.g., i.e. hacking, cracking and attacks). So while hackers stealing credentials will are considered, fires at a particular site's data center are not.

Resources

The initial risk assessment took 1 year to complete and required a whole team with representation from every major operations group led by a security analyst. Updates every couple of years require much less time and resources, typically a security analyst dedicating a fraction of their time over a few months and support from the XSEDE Security Working Group.

Approach

The NIST 800-30 risk assessment approach was chosen for several reasons as described in 02: A Risk-based Security Program for XSEDE. The risk assessment updates will proceed in the same order through the same 10 steps as the original, with the security analyst updating each section, waiting for feedback from the XSEDE Security Working Group, incorporating feedback, and then moving on to the next step. At the end of the update process, the risk assessment report will be shared with senior management.

04: XSEDE System Characterization

(2019 Note: Some links in this document may not be active due to their target in the old XSEDE Wiki which has been replaced by https://confluence.xsede.org. Several Wiki links do not have corresponding Confluence links.)

The first step of any risk assessment is to concretely define the boundaries of the system(s) that you are evaluating. This keeps later steps, such as identifying threats and vulnerabilities, much more focused. It also provides a systematic way to break up the evaluation of threats and vulnerabilities by logical areas that can be delegated to the most appropriate person(s).

This is a federation-wide risk assessment, which may in the future need to reference individual site-level risk assessments for more detail. We, however, do not have the expertise or resources to evaluate risks to individual compute resources at the different service providers.

That said, the purpose of this phase is to identify the boundaries of the "system", along with the resources, organizations and information that constitute the system. This includes identifying the system's mission; major hardware components; key software and services; data and information with sensitivity assessment; user and support communities (includes XSEDE staffing groups as well as resource providers); logical network topology; system interfaces (both internal and external connectivity); flow of information (particularly if sensitive); references to existing security architecture and policy documentation; and management and monitoring controls (e.g., availability monitoring, intrusion detection systems).

System Mission

XSEDE's mission is to substantially enhance the productivity of a growing community of researchers, engineers, and scholars through access to advanced digital services that support open research.

Users

Broadly speaking, our users are principal investigators (PIs) or students/researchers working for PIs. PIs are research scientists or educators at a U.S. institution. Most, but not all PIs, are current grant holders from the NSF.

Besides allocations given to XSEDE staff for testing and development purposes, there are three basic types of allocations: startup, education and research. Startup allocations are smaller allocations that are primarily used by first time XSEDE users and are not hard to get. Many universities have XSEDE campus champions who can help people get start-up allocations. Education allocations are for classes or training at workshops and such. The primary and main use of the systems is for research allocations whose requests are reviewed at quarterly XRAC meetings. These are longer-lived and higher priority as this is where the science is done that fulfills the primary mission of XSEDE.

Users have many issues and concerns, though the primary security concerns will be availability and integrity of their data and results. If there is doubt about the integrity of their data or results, the work is useless, and hence that is the primary concern. However, resources that are unavailable due to outages or security incidents are nearly as problematic as many jobs require long runs without interruption. Concerns about confidentiality likely come in third for most users as XSEDE is used for fundamental research which is generally open, and not for private sector uses that have more proprietary concerns. An institution may provide an HPC that serves both, but it is unlikely that the private sector customer is an actual XSEDE user using an XSEDE allocation and is thus tangent to this risk assessment.

See https://www.xsede.org/allocations for more information.

XSEDE Staff

For this subsection, please refer to the organizational chart below (current Dec. 2011).

XSEDE is led by a project director who receives input and direction from the several stakeholders: the National Science Foundation (NSF), the XSEDE Advisory Board (XAB), the User Advisory Committee, the XD Service Providers Forum, and the XSEDE Project office (which the director leads). The project office members cover 6 areas: External Relations (aka, public affairs), Industry Relations, Systems & Software, Architecture & Design, Project Management, and Software Development & Integration.

Besides the project office, there are 5 additional areas led by level 2 managers. These are Operations, User Services, Advanced User Support Projects, Advanced User Support Communities, and Education & Outreach. The XSEDE Senior Management Team is led by the project director and consists of the level 2 managers plus some additional members of the project office. It currently includes the directors of Operations, Users Services, Extended Collaborative Support Services-Projects, Extended Collaborative Support Services-Communities, Education and Outreach, the Senior Project Manager, the Senior Systems Engineer, and the Software Development and Integration Lead.

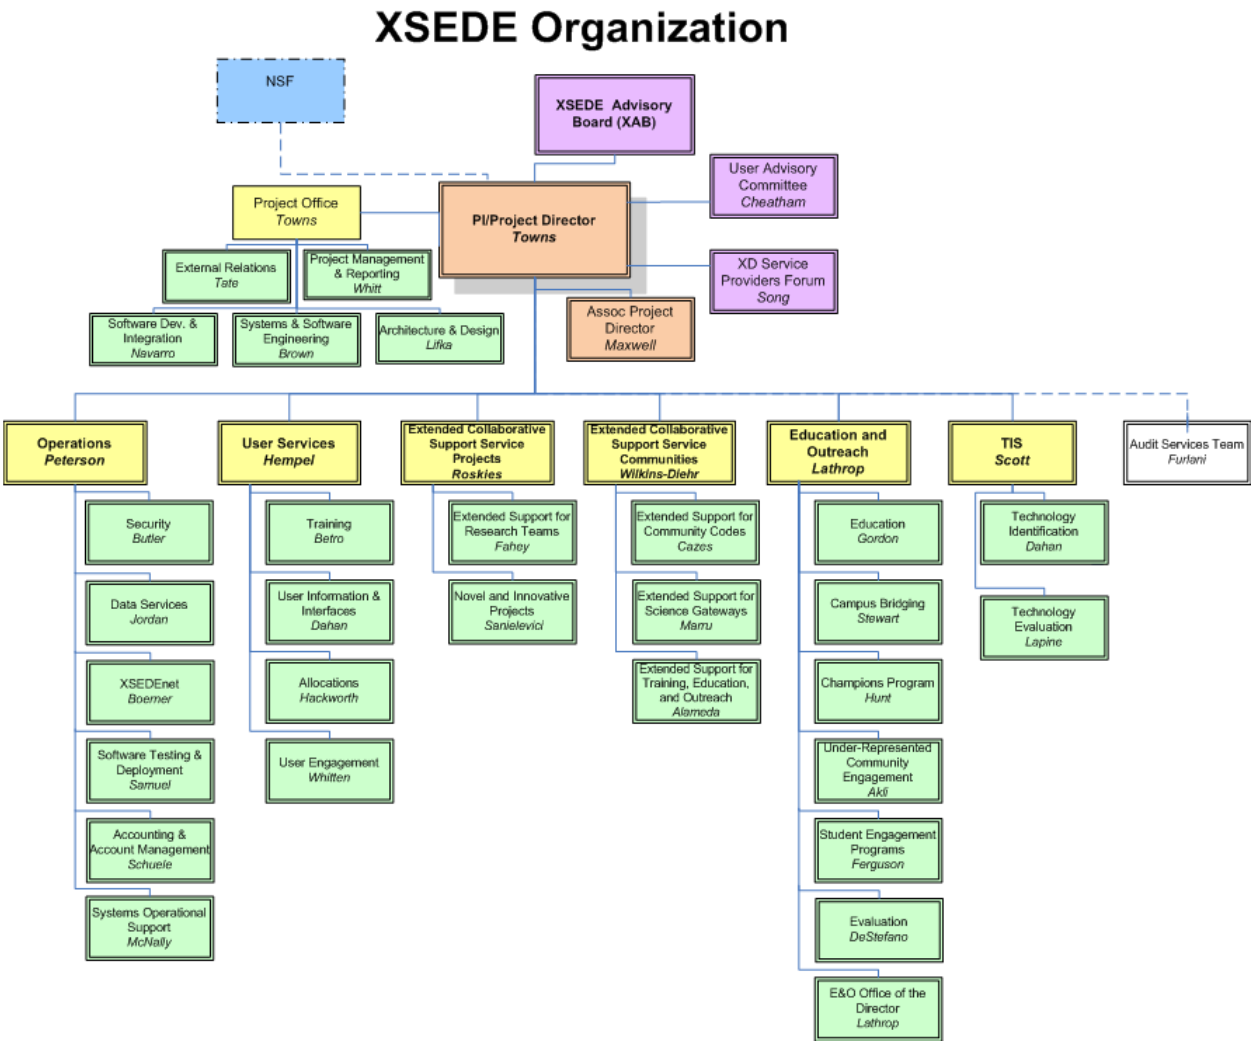

From Concept to Deployment

The parts of the organization that affect the design, deployment and operation of XSEDE systems and services are of primary concern to this system security risk assessment, though clearly failures in other divisions can introduce project risks, if not cybersecurity risks. For example, the failure of the education and outreach program or public affairs could easily derail future refunding efforts.

The teams responsible for design, deployment and operation of systems are the Systems & Software Engineering Team (SSE), the Architecture & Design Team (AD), the Software Development and Integration Team (SDI) and the Operations group. The general process followed is that SSE defines the requirements, holds the baselines, and interfaces with the advisory committees to gather the information they need. They also are responsible for ensuring that XSEDE uses best system and software engineering practices across the full life-cycle of engineering effort. Requirements and baselines are fed into the design process which is then carried out by the AD team, hopefully, with back and forth feedback between SSE and the AD teams. SDI is then responsible for detailed design, development, and integrated testing. Anything new must fit into an existing ecosystem, and SDI is supposed to get any new services ready for development with an integration plan for Operations. Finally, Operations works with SDI to deploy new technologies and runs them on a day-to-day basis. Most actual software development occurs in SDI, though lines are in practice a little blurrier than described above.

Staff Demographics

Staff for XSEDE are professionals and academics primarily located research universities or DOE labs with HPC resources. There is a strong correlation between the major XSEDE service providers and the location of staff, and hence staff are spread out quite a bit geographically (primarily at NICS, TACC, NCSA, UChicago, IU, and PSC). Staff include developers, system admins, project managers, post docs, faculty and more.

Because of geographic disparity, most communication is in person at quarterly meetings or on the many conference calls and emails. The fact that staff is so spread and so numerous is important to note, as it does lend itself more readily to miscommunication and social engineering attacks.

Service Providers

The relationship between XSEDE and various partners range from tightly-coupled (e.g., those providing central services and XRAC allocated compute resources), to loosely coupled (e.g. providers of a limited set of XSEDE compatible services), to peripheral (e.g. entities only wanting to make the services they provide visible to the XSEDE community). As such, Service Providers (SPs) are categorized into a tiered system with representation on the Service Providers Forum (SPF) which in turn will have representation on the XSEDE Advisory Board (XAB). How much representation an SP has in the SPF is determined by which of the three tiers it resides in, which in turn is defined by the minimum requirements they meet and agree to in formal agreements with XSEDE.

SPs classified as Level 1 Service Providers have the deepest level of commitment and integration with XSEDE and the XSEDE environment, and they explicitly share digital services with the broader community of users of the XSEDE environment and infrastructure making use of XSEDE services and interfaces. Level 2 Service Providers make one or more digital services accessible via XSEDE services and interfaces, and share one or more digital services with the XSEDE community along with the organization's local users. Level 3 Service Providers are the most loosely coupled; they will advertise the characteristics of one or more digital services via XSEDE mechanisms, might make those resources or services accessible via XSEDE compatible interfaces, and are not required to share them with the XSEDE user community.

Level 1 SPs are tightly integrated and provide some of the core cyber-infrastructure for XSEDE in addition to the major HPC resources. Certainly, those providing cyber-infrastructure, such as shared filesystems and Authentication/Authorization services, should perform their own risk assessments for a complete understanding of XSEDE risks. However, even those not providing shared cyber-infrastructure but just HPC resources should also do risk assessments since (1) they are automatically trusted more being on XSEDENet, (2) tight integration means their failure could have unexpected system-wide implications, and (3) failure of key resources impairs the mission of XSEDE, and (4) the public often does not differentiate these sites from XSEDE which means a security incident at any major SP affects the reputation of XSEDE as a whole.

Level 2 resources are not running key resources or core cyber-infrastructure, but they often do have some special integration with XSEDE. Furthermore, they will often install pieces of the XSEDE software stack and thus depend on XSEDE's security to some extent. XSEDE does not gain much additional risk from these partners, but they may gain a little risk by integrating into XSEDE. However, it is hard to make such blanket statements as each of the level 2 agreements and integration points are likely to be unique. Therefore, while it is not practical or necessary for each level 2 resource to perform a risk assessment of their own for submission, security operations in XSEDE must be aware of these integrations ahead of time and part of the approval process for connecting new resources.

Level 3 SPs are expected to be mostly just listed through the XSEDE portal or some other directory, while operating quite independently. They pose little to no risk to XSEDE. If they do make resources available through XSEDE interfaces, that mechanism should be understood and approved by the XSEDE security operations team, but there is no need for individual risk assessments of these SPs.

Based on current (Feb. 2014) information, we have the following level 1 SPs, none of which have done individual risk assessments for XSEDE.

- NICS (National Institute for Computational Sciences)
- PSC (Pittsburgh Supercomputing Center)
- SDSC (San Diego Supercomputing Center)
- TACC (Texas Advanced Computing Center)

Systems & Services

Each major type of system or support service for XSEDE is listed in this section with links to where we describe (1) the data and information on the system, (2) the flow of any of that data between systems, and (3) the relevant interfaces to the system or service. In particular, with the data we want to identify any sensitive data and whether or not it cross organizational or trust boundaries. System interfaces could be either internal backend interfaces where credentials are passed, or more external interfaces like a shell access exposed through a login node. For any identified data flows, there must be some corresponding system interface(s) listed.

Details about these services can be found in Appendix A.

SERVICE	NOTES
Level 1 SP Compute and Storage Resources	HPCs and archives
AMIE	Backend for XDCDB
CI Logon	
CI Tutor	www.citutor.org
Cornell Virtual Workshop	www.cac.cornell.edu

XSEDE.org DNS	
Conference Call System	Lync
E-Mail List Server	
Genesis II	STS & RNS servers
Globus Listener	globus-usage.teragrid.org
Globus Online	www.globusonline.org
Metrics for IIS	info-metrics.xsede.org
Integrated Information Services	info.xsede.org
INCA	inca.xsede.org
Jira	jira.xsede.org
Karnak Predictor	karnak.xsede.org
XSEDE Kerberos Realm	
MyProxy Certificate Authorities	myproxy.xsede.org
Nagios Service	nagios.xsede.org
OAuth	OAuth redirect on portal site
Openfire Secure Jabber Server	For incident response communication
perfSONAR	Network monitoring
Resource Description Repository	rdr.xsede.org
RSA SecurID	
RT Ticket System	tickets.xsede.org
Sciforma	projects.xsede.org
Security Wiki	For incident response
SharePoint	Reporting wiki
Single SignOn Hub	login.xsede.org
Source Repository	software.xsede.org
Speedpage	speedpage.psc.edu
User Profile Service	
XDCDB Central Database	xdcdb.xsede.org
XRAS Service	Proposal system
XSEDE User Portal, KnowledgeBase, Web Site and Wiki	www.xsede.org

An updated list of these services is maintained at https://sysops.xsede.org/xes-index/ .

Network Topology

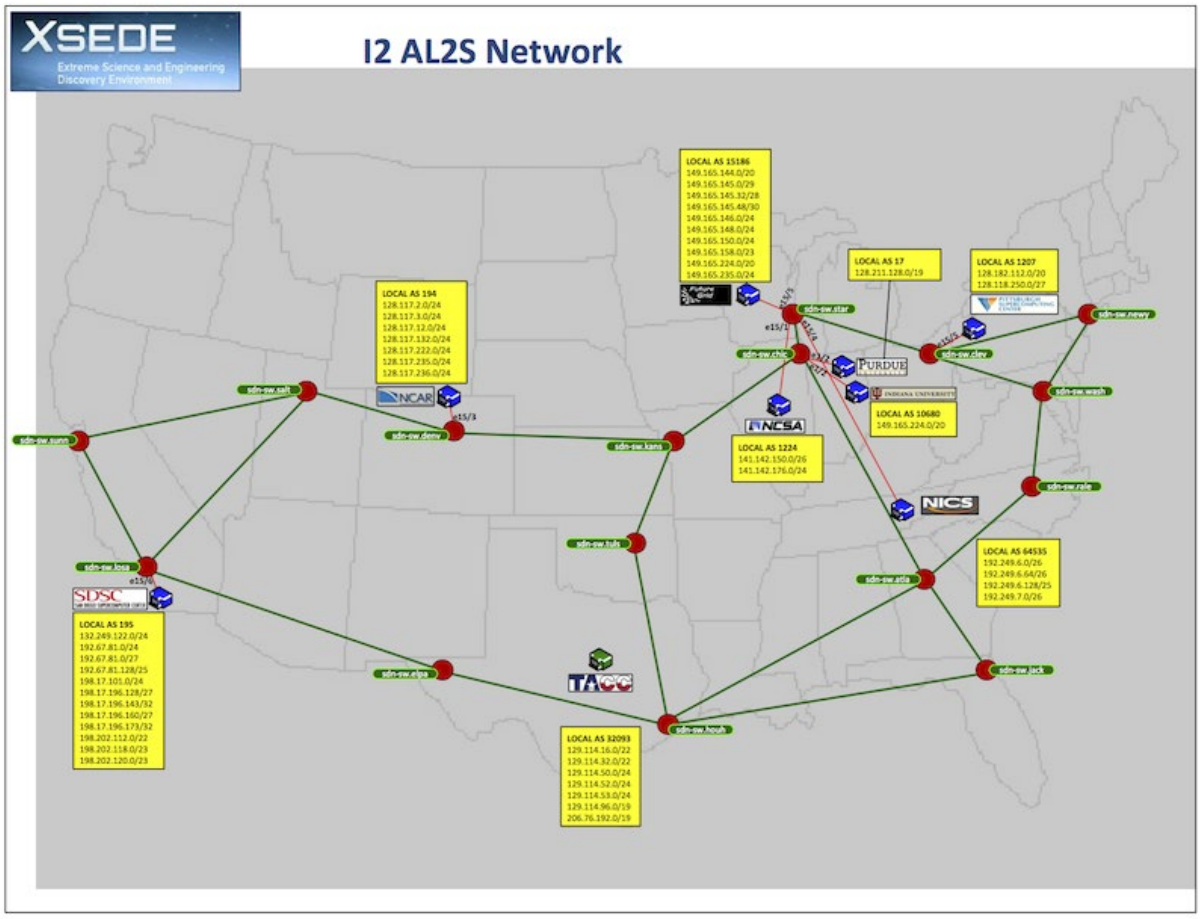

The diagram above shows the 10,000 ft view of XSEDEnet, a high-speed, wide-area backplane used for perfsonar, gridftp and other large transfers between large XSEDE SPs. Each site uses its own set of 10G fiber links to connect to their local POP< and the backbone is provided through Internet2.

XSEDEnet does not own any of its own network equipment, rather this is all owned by our service provider, Internet2. Of course, the SPs own the router endpoints in their own data centers. XSEDEnet also has a dedicated /25 IP space donated by NCSA, though each site uses their own IP space to advertise their resources into XSEDEnet as needed. SPs are only supposed to advertise routes for addresses of XSEDE specific resources on XSEDEnet, and the other XSEDE sites prefer routes over XSEDEnet to those resources. As such, there is some implicit trust and most sites do not run any firewalls or perform any extra monitoring on XSEDEnet. However, each site can make its own decision here, and some sites (e.g., NCSA) will monitor the XSEDEnet like any other WAN connections. Therefore, the risk from compromise of XSEDEnet is better evaluated one level lower at the individual SPs.

XSEDEnet is also not critical to XSEDE in that sites should fail-over to other research and education network links (at a cost to performance possibly). However, that has not been tested in practice.

Finally, it should be noted that XSEDEnet is meant primarily for Level 1 SPs. Level 3 partners won't use XSEDEnet, and some level 2 SPs do not use it.

Management & Monitoring Controls

Management structure and lines of authority for XSEDE are described above in the "XSEDE Staff" section. Procedural controls related to security are described in the incident handling playbook and other documents below under "Existing Standards and Policies". Environmental controls, contingency and recovery controls, personnel and physical security controls are largely absent or not applicable to XSEDE.

For each particular system and service, we indicate management interfaces available to XSEDE staff in the "System Interfaces" sub sections. However, many of these services are monitoring controls themselves at a higher-level. These monitor the state and availability of critical resources such as HPC systems and network links and include: perfSONAR, IIS, Metrics for IIS, INCA, Speedpage, Globus Listener, Karnak Predictor, and the User Profile Service.

XSEDE Software Stack

This section identifies key pieces of software, such as toolkits and services, run at different service providers (SPs). What isn't included are local customizations, particulars of base OS'es and other software run at the individual sites that would need to be addressed in individual, site-specific risk assessments. Instead we are focused on software (e.g., Globus Toolkit, Unicore, etc) and configurations (e.g., the XSEDE trust store of certificates and Common User Environment) that are configured in XSEDE specific ways to support the XSEDE services layer in the XSEDE architecture. Much of this content is derived from the XSEDE Production Baseline for Service Provider Software and Services.

SOFTWARE/SERVICE NAME	REQURED ?	DESCRIPTION
Registration Services	Partial	The Core Integration Capability Kit must be installed on all SPs and includes tools (pacman, and ctss-core-registration) needed to publish information back to Integrated Information Services regarding the capabilities of each SP, i.e. what resources they provide. Additional registration subcomponents that are required include data movement client/server registration, local compute registration, login service registration, and visualization tool registration. Other registration services are optional. This includes application runtime and development support, distributed programming systems, local hpc software, meta-scheduling registration, co-scheduling, parallel application, remote compute (job submission), science gateway, DC-WAN Lustre, and workflow system registration.
INCA	Yes	INCA client to allow verification of registered services at the SP.
AMIE	Yes	AMIE client to report back accounting information to the central database

Globus Toolkit	Partial	Globus Toolkit is a suite of tools provided by the Globus team at http://globus.org/toolkit/. Not all of these tools are required for XSEDE SPs, but GridFTP servers for file transfer, globus-url-copy (the corresponding gridftp client tool), GSI OpenSSH (which supports X.509 certificates), and a MyProxy client (to request X.509 certificates) are required. Other Globus Toolkit client components are optional, such as, and GRAM
XSEDE CA Tarball	Yes	There is a set of certificate authorities which all XSEDE SPs are supposed to install and trust for GridFTP and GSI OpenSSH interoperability. These are maintained in a tarball in a subversion repository at software.xsede.org.
UberFTP	Yes	UberFTP is a gridFTP client developed at NCSA.
Globus-wsrf	Yes	This refers to the web services resource framework container for the Globus services such as GRAM. This does not mean that the services themselves must be installed at the SPs
XSEDE Glue2 Publishing	Yes	The XSEDE GLUE2 component is a module that publishes GLUE 2.0 specified information thru information services at each SP. We aggregate all the SP GLUE 2.0 information into central information services for use by the Condor meta-scheduler and other software components.
Local Resource Management System	Partial	This refers to the various local job schedulers (e.g., Cobalt, Condor, LSF, Load Leveler, Moab, PBS Pro, PBS Torque, and SGE). Supported schedulers should be registered with ctss-local-compute-registration. Not each and every one of these workload management services need to be available at every SP.
gx-map	Yes	A tool for requesting updates to the gridmap file to associate new distinguished names with a user account name.
xdusage	Yes	This is a command line tool that allows users to see the status of their current allocations.
Modules	Yes	Modules is the XSEDE standard environment management tool. It is often used to manage multiple versions of software and supports most shells and some scripting languages. See http://modules.sourceforge.net/ for more information.
tgproxy	Yes	This software manages the location and naming conventions of user proxy credentials.
Common User Environment	Yes	The Common User Environment (CUE) is a standardization of user environment variables aimed at providing a more homogenous user experience across XSEDE SPs.
tginfo	No	This is a command line tool to query the XSEDE Integrated Information Services.
tgresid	No	Site, resource, and platform type identification tool.
SAGA	No	A library for grid application programming.
MCP	No	The Master Control Program (MCP) optimizes application start times by submitting multiple copies of an application to different resources. Once one copy begins to execute, the other copies are deleted.

GUR	No	The GUR (Grid Universal Remote) tool is a python script which uses the ssh, scp, or their gsi counterpart commands to help users make reservations, compile programs, and co-schedule jobs. It can accommodate flexible node range and start/end time.
CommSH	No	The Community Shell tool is a restricted shell used for community accounts for science gateways. Since such users have a defined set of allowable actions, a restricted shell can be used. This mitigates the risks of having many users share a single account.
Unicore	Yes	This software provides the basic execution service to XSEDE compute resources with the gateway, Unicore X and TSI software components.
Genesis II	No	In XSEDE Genesis II clients can utilize the Unicore basic execution service (BES). This depends on Genesis II root name service and secure token service described above. GFFS is not currently deployed in XSEDE nor the Genesis II native BES.

Teragrid Policies and Procedures

XSEDE inherited policies from TeraGrid, and they need to be updated. Of particular interest to security are TG-4, TG-5 and TG-10, which cover incident response, access control and account policy, respectively. Security baselines and security standards can be found at http://security.teragrid.org/. XSEDE does have a User Responsibility Form, which is very similar to an Acceptable Use Policy (AUP) that most IT service providers have.

New Policies and Agreements

In addition to the aforementioned policies that should be updated, there are some others that do not exist but maybe should.

- A federation charter which establishes governance structure, partner rights and mechanisms to resolve disputes.
- A legal, law enforcement and regulatory plan. Bits of this are spelled out in the incident response handbook, but the topic is far from completely covered.
- SP security responsibilities. Any official agreement with SPs should have formal language about security relevant responsibilities, just as an AUP for users would.
- PI agreement. The PIs take on an extra responsibility for their allocations above and behind that of an ordinary user.

Security Architectures

There is no official security architecture for XSEDE. However, there is a security production baseline for authentication and authorization.

Password Policies

XSEDE uses x509 certificates issued by trusted CAs to allow logins, job submission and file transfer. There are multiple ways to authenticate to these trusted CAs and get certificates that can be used on a users' behalf. Because these certificates can be used at any level 1 SP, the strength of the whole authentication system is only as strong as its weakest length. Therefore, we survey the password requirements of each entry point to gain a certificate, including all the level 1 SPs, the XSEDE User Portal (XUP), InCommon which can be linked to XUP accounts, and the science gateways which create proxy certificates for users.

SITE	LENGTH	COMPLEXITY	COMMENTS
XUP	>=8	3 of following: upper, lower, numbers & symbols	Uses XD Kerberos
NCSA	>=8	3 of following: upper, lower, numbers & symbols	
NICS	10-12	Digits only	Uses RSA One Time Passwords
PSC	>=8	3 of following: upper, lower, numbers & symbols	
SDSC	08-08	Must pass cracklib check	Not IGTF accredited
NERSC	>=8	Combination of upper and lowercase letters, numbers, and at least one special character within the first seven positions, and nonnumeric letter or symbol in the first and last positions.	
InCommon	none	none	Defined by Universities (NO STANDARDS)
Science Gateway (traditional)	none	none	Defined by gateway (NO STANDARDS)
Science Gateway (OAuth)	N/A	N/A	Same as XUP as account needed there
DOEGrids	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password
ESNet	N/A	N/A	Does not directly give out certificates to users but certifies other CAs like DOEGrids
IRISGrid	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password
UK E-science	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password

NIKHEF	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password
KEK	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password
DigiCert	N/A	N/A	JIM BASNEY PLZ FILL IN
INFN	N/A	N/A	User must visit and RA in person, cannot get a certificate just based on password

Publickey Auth

Public key authentication is not allowed by all SPs, unlike GSI x.509 based authentication. Therefore, risks of exposed or shared keys are best addressed in SP-level risk assessments.

Process & Examples

- NIST definition of system characterization and example survey
- Short 1-page table system characterization template
- Common example risk assessment. It is too low-level for us, but it is also a much simpler system.
- Another example. Better level of detail for us, but still a much simpler system.

05: Threat Profile

Purpose and Goals

This threat profile discusses likely threat sources that may attempt to exercise vulnerabilities of the XSEDE trust fabric or shared resources, whether intentional or not. Threat sources considered here are humans (insider/outsider threats) or software/application errors that affect more than one XSEDE service provider or core shared services. Other threat sources, like natural disasters, are too site specific and should be addressed by risk assessments done at the major service providers individually.

In addition to identifying potential threat sources, an effort is made to estimate the motivations, resources and capabilities of such threat sources. We focus upon threat sources that would likely target an organization like XSEDE and list threat actions they may take to exercise potential vulnerabilities. We specifically note the threat sources not considered here and why they are not.

The process followed comes from the NIST 800-30 standard for IT security risk assessments. The relevant portions of that standard and other threat profiles used as templates are found in the Sources Utilized section below. Further, the results of the surveys sent to the tier 1 XSEDE service providers are found in Appendix D. These surveys questioned security teams at those organizations about past Teragrid relevant security incidents as well as potentially new threats with XSEDE. The information from those surveys informed the table in Appendix C. This table maps threat sources with potential motivation with potential actions against XSEDE entities, and it informs much of the text of this document. It also tries to discern already seen or expected threats from those that are unlikely but could possibly exploit XSEDE vulnerabilities.

The results of the surveys and other content may be sensitive, and therefore must not be exported with this page and shared externally. Nothing directly contained within this threat profile is sensitive, but permission must be sought from the XSEDE CSO before sharing it externally.

Human Cyber Threats

From experience in Teragrid and the first 4 years of XSEDE, along with the consideration that our user's data is for open scientific research, we expect three kinds of human threat sources in XSEDE: crackers, computer criminals, and insiders.

Crackers

Crackers, often referred to as hackers, come in a range of skill-levels with varying motivations. The most common, and least worrisome, is the script-kiddie, who has a very low skill-level. Often, they do not even know how to program. By virtue of being on the Internet, we are exposed to constant script-kiddie attacks, such as, port scanning and dictionary-based SSH attacks against hosts. However, not all crackers are so unskilled as FBI Major Incident 216 showed the Teragrid community in 2004. A very skilled and determined teenager from Sweden appears to have single-handedly cost thousands of man-hours of incident response time across many sites with his constant attacks over the span of 14 months. Several sites had to go offline or reset all user passwords.

Motivation is hard to determine for any particular incident, especially if it does not result in an arrest. Most often crackers are motivated by ego, the challenge, rebellion or some form of activism. We have not seen much if any activism motivating attacks in XSEDE or Teragrid. Activists usually deface sites or tell you why they hacked you. They want publicity. While some survey respondents noted the rise in hacktivism elsewhere with LulzSec, Anonymous and others, we have not seen it in our community yet. Furthermore, XSEDE as service provider of open scientific research resources just doesn't attract hacktivism like a

federal lab or .gov site might. Unless XSEDE resources start being used for more politically sensitive work (perhaps climate change modeling), it seems unlikely that this risk will change significantly.

The most commonly seen problems from crackers, and ones likely to continue, are dictionary-based and credential harvesting attacks. Dictionary-based attacks that attempt thousands of passwords are noisy and only rarely successful. They typically succeed where services are poorly configured, and they usually represent the actions of script kiddies. The more advanced attacker has the more common MO that we see in the most common type of incident. The pattern is as follows. Credentials are harvested, often at a non-XSEDE resource. This password is then used to login to an XSEDE resource where upon a local privilege escalation attack is perpetrated. This may mean waiting for the next Linux zero-day exploit, at which point we see root escalation from potentially multiple accounts. The next step is often to trojan another service to either keep a foothold or start harvesting more credentials. Then the cycle repeats. This particular threat is unlikely to end or abate any time soon unless password and publickey authentication is removed as an acceptable SSH login method, or interactive shell access is no longer available to users.

There are several other common threat actions that crackers may take but that are less likely. They may deface web sites, perform cross-site scripting to try to install malware on user workstations, socially engineer our accounting or administrative processes, use our hosts as proxies to hide behind, or use our systems as a foothold to monitor the private and trusted XDNet. None of these actions were reported as likely by survey respondents. It has been noted that HPCs can be particularly useful at cracking passwords, and in the distant past at least one was used to do so. However, this has not been seen in a very long time, probably because passwords are much easier to crack on commodity devices now, and few outsiders actually know anything about submitting HPC jobs or writing code for them.

Computer Criminals

While the term computer criminal is not completely mutually exclusive of the label cracker, we attempt to make a distinction here because their motivations and resources are different enough. Furthermore, the types of threat actions they may take differ enough to justify separate discussion. Being motivated primarily by direct monetary gain or perhaps use of a high-powered attack platform, the actions they take may differ. Also, criminal organizations often hire many skilled programmers for their maleficence. This could mean that they have many more resources to throw into an attack if it is directly targeted. Fortunately, there is a lot of low-hanging-fruit and it is not usually in their interest to focus on a single adversary. Their strategy is usually breadth-first, gathering as many small hosts as possible from many locations.

Cyber criminals most commonly are found "bot-herding", where they make large untargeted attacks to gain as many hosts as possible, which they collect as a commodity resource. They may do this through attacks seeking to exploit remote vulnerabilities or by dictionary-based password guessing against services like SSHD. As with crackers, the main defense against this threat source is a well-maintained and well-configured system. Cyber-criminals can also follow the more involved steps described earlier for crackers, that is, harvest credentials, perform a local privilege escalation, and install a trojan, backdoor or malware. The difference here is that the malware they install is often for activities like spamming, hosting illegal content or performing DDoS (distributed Denial-of-Service) attacks.

Just as with crackers, they could perform social engineering or simply be looking for proxies to hide behind, though this has not been a noted activity of cyber criminals in XSEDE or Teragrid to date. Just like crackers, these adversaries typically aren't very aware of HPC environments and are not specifically targeting HPCs for their potential. That has been the case to date, but some other HPC sites have seen a targeted type of activity whereby attackers would submit hidden jobs that perform bitcoin mining (a way to earn virtual, anonymous currency by performing computations) on HPCs (This happened in XSEDE recently, but by an insider and not a traditional cyber criminal). These were targeted attacks by persons with more domain specific knowledge than most. This new revenue stream does change the economics and increase the risk of future targeted attacks against XSEDE resources by cyber criminals. While it has

been supposed in the past that these powerful machines with high bandwidth could be targeted by criminals to use as DoS tools, this has never really been realized as a threat. It is probably because DDoS attacks on commodity machines are easy enough to gather, just as effective, and harder to block at the network level.

Insiders

Insiders are either XSEDE employees or users who already possess valid credentials to access XSEDE resources. Someone could be motivated to do harm if they are a disgruntled or former employee. However, no one has provided an example of that happening in Teragrid or XSEDE. An insider could also violate policy or exceed their authorizations simply because they are motivated by curiosity about another's research or personal information. Examples of this are rare but have happened. Most commonly "insiders" aren't motivated per se but are responsible for unintentional errors which still cause threats to system stability and security. Someone could simply be negligent, poorly trained or have made a programming or configuration error which affects the availability, confidentiality or integrity of an XSEDE system. The exception to this was the purposeful mining of bitcoins by a staff member and PI. This did not harm the system but was an intentional action that misused resources for personal gain.

The most common threat actions seen by insiders in the past were unintentional system bugs or misconfigurations, which have overly exposed data or negatively affected availability of a resource. Some users have been found trying to browse another user's data, though rarely. Such threats could be realized through similar privilege escalation and trojaning attacks seen from other threat sources. Those surveyed have not seen or expected system sabotage or data corruption/deletion, though that is always a potential threat from insiders. Other possible but unlikely insider threats include blackmail, malware inserted to infect other users, or interception and monitoring of local network traffic.

Software & Configuration Threats

Software errors and configuration threats were touched upon under the category of human insider threats, but they could also be categorized by themselves separately. They are a different enough kind of threat that they probably merit a separate discussion. While such threats sources often have an insider responsible, such as a system administrator or programmer for XSEDE, it need not be tied to an insider, and hence it becomes somewhat detached from the notion of a human actor. For example, XSEDE may depend on software that it does not develop, and the quality and support of such software may very much impact the stability and security of XSEDE. If a key piece of software has an exploit against it in the wild, XSEDE Operations staff need to be sure that the party responsible will quickly respond with appropriate patches or remedies. More indirectly, software without an exploit that is just unreliable reduces the reliability and stability of XSEDE, and hence it becomes a security concern as availability of resources is a pillar of security. Finally, it is possible that software used is simply out dated or completely unsupported. There is no real human actor behind such a threat source, unless you want to say that the threat source is an inept system administrator. However, that may not be fair as we may simply be stuck with the software for historical reasons. But such risks still need to be identified and mitigated if possible, and issues with major software libraries like OpenSSL have only made this issue more relevant in recent years.

Threats & Threat Sources Out of Scope

As we say in 03: Risk Assessment Project Sizing, this is a narrowly scoped security risk assessment. We are only considering insider/outsider threats (e.g., hacking, cracking, attacks, etc.) and software/application errors. We do not address physical or environmental threats. This is because we are doing a federation-level risk assessment focused on threats that affect core/shared services and more than one service provider. Natural disasters, physical threats and threats specific to individual HPCs

would all be best addressed one layer lower with risk assessments done by the service providers themselves more directly.

There are some other threats we don't consider simply because we do not have the assets that would be targeted. In particular, industrial espionage is not a threat considered here because XSEDE is not using or producing commercial intellectual property. This work is geared towards open, scientific research for the NSF. While some sites may sell cycles to industry, and thus have to worry about industrial espionage, it is not an XSEDE specific threat. The only effect on XSEDE is the unavailability of specific resources shared by XSEDE and industry partners at specific sites, and the impact of that threat on the availability of resources is site-specific and best determined at lower-level per site risk assessments.

Finally, there are some threats that are simply not possible to address in the open environment of XSEDE or with the resources it has. For example, a directed attack from a nation state actor would be such an asymmetric threat that we could not protect against it. If we are lucky, we might be able to detect such attacks, but even that is questionable. Procedures to protect both technically and socially against nation states would be far too draconian even if we did have the resources. Furthermore, there is little gain for nation states or terrorists to attack XSEDE. We have no critical infrastructure and no loss of life or major monetary damage could be brought against the United States by bringing XSEDE offline.

Summary Statement

The threats to XSEDE largely have not changed from the threats to Teragrid, besides some increased risk of credential harvesting with many additional tier 3 resources expected to join. Therefore, our historical data can intelligently inform our threat profile going forward. We expect the most common problem to continue to be credential harvesting, followed by privilege escalation, then trojaning, and finally more credential harvesting. This threat is unlikely to change without major changes to authentication mechanisms and/or user interfaces. We will continue to see many, mostly unsuccessful dictionary-based attacks as well as bot-herding for criminal activities. We anticipated and saw a new threat action of bitcoin-mining. Other than that, our biggest threats are misconfigurations or software errors that must be addressed through configuration management and good software engineering practices. Fortunately, we avoid some of the threats seen against other sites by the nature of the customers we serve. However, if that ever changes and we start serving more private industry with valuable intellectual property, governments with classified data or researchers working on politically hot topics, we will have to reevaluate our threat profile.

Sources Utilized

- Threat Profile Process from NIST
- Blue Waters threat profile
- Early Phase Blue Waters threat assessment

Threat Profile Survey

We used surveys sent to the tier 1 XSEDE SPs to gather historical data. The survey questions and responses are available in Appendix D.

06: Vulnerability Identification

A vulnerability is a flaw or weakness in system security procedures, design, implementation or internal controls that could potentially be exercised by a threat agent to result in a breach or violation of the system's security policy. In this step of the risk assessment we identify as many federation-level vulnerabilities in XSEDE as we can, mapping each vulnerability to a threat action and threat source where possible. If there is no corresponding threat, we can ignore the vulnerability in further stages of the risk assessment.

Survey

Questions are asked of the XSEDE Security Working Group members and representatives of each area in XSEDE operations. The purpose is to identify as many security vulnerabilities in XSEDE as possible, and the corresponding threats which could exercise these vulnerabilities (See examples in the table below). In future steps, we will add probability and impact to each item in the table, thus creating a set of XSEDE risks. However, this step is about completeness, and we are trying to enumerate all vulnerabilities that we can, regardless of the likelihood or impact.

Vulnerability Identification survey questions and answers can be found in Appendix E.

Example Vulnerabilities

VULNERABILITY	THREAT SOURCE	THREAT ACTION
Terminated employees' system identifiers (ID) are not removed from the system	Terminated employees	Dialing into the company's network and accessing company proprietary data
Company firewall allows inbound telnet, and guest ID is enabled on XYZ server	Unauthorized users (e.g., hackers, terminated employees, computer criminals, terrorists)	Using telnet to XYZ server and browsing system files with the guest ID
The vendor has identified flaws in the security design of the system; however, new patches have not been applied to the system	Unauthorized users (e.g., hackers, disgruntled employees, computer criminals, terrorists)	Obtaining unauthorized access to sensitive system files based on known system vulnerabilities

Data center uses water sprinklers to suppress fire; tarpaulins to protect hardware and equipment from water damage are not in place	Fire, negligent persons	Water sprinklers being turned on in the data center

Identified Vulnerabilities

Below is a table of all vulnerabilities known or ones that can be reasonably inferred (e.g., there may be no known user portal vulnerability, but it is safe to assume any complex web site can be compromised). Information was gathered by investigating assets from 04: XSEDE System Characterization, reviewing other risk assessments, interviews with people from each operations domain, and review and acceptance by the risk assessment team. Each row of this table contains a vulnerability, a threat source taken from 05: Threat Profile, and an actual threat that would exercise the vulnerability. As Shon Harris states, "a risk is the likelihood of a threat agent taking advantage of a vulnerability and the corresponding business impact." Therefore, each of these numbered rows were mapped to at least one risk as we added likelihood and impact analysis in future steps.

It is worth noting that XSEDE has an existing Risk Register. This was created separate from this exercise and not by a formal process. Furthermore, its scope is much more general. Once this security risk assessment was approved, its unmitigated risks were added to the XSEDE risk register.

VULNERABILITY	THREAT SOURCE	THREAT ACTION	LIKELIHOOD	IMPACT	TOTAL RISK	INDEX
Access Control						
Shared accounts making auditing difficult	Malicious user or Cracker	Science Gateway accounts are shared. Incident response involving gateway accounts may require the assistance of a third party if appropriate attributes are not passed along, thus allowing attackers to better obfuscate their identity.	L	L	1 Low	1
Inconsistent authentication standards make XSEDE authentication only as strong as the weakest SP.	Crackers	Through use of x.509 certificates, users have single-sign-on capabilities to move between level 1 SP resources. The requirements for authentication mechanisms, identity vetting and operating a CA between SPs are different, making the whole system only as strong as the weakest link. However, all the CAs are IGTF accredited now, and the password policies are roughly equivalent among the XUP and level 1 SPs. So while an attacker might	L	L	1 Low	2

		leverage a slightly weaker password standard at one site, the bigger risk is from reliance on passwords which users often reuse.				
Some XSEDE services use their own authentication system which is not subject to the same security requirements or account deactivation processes.	Disgruntled former User/employee or Cracker	A few services (e.g., Sharepoint and CI Tutor) use one-off authentication systems separate from the XSEDE kerberos and CA systems. Users could exploit the fact that disabling an XSEDE kerberos account does not deactivate these, and adversaries could exploit unknown vulnerabilities in these outside authentication systems if users synchronize their passwords for the portal and these.	L	L	1 Low	3
Anyone can create an XSEDE account and identity vetting is dependent upon trusting email, delegations to PIs and self-asserted attributes.	Cracker or malicious user	Having a lower level of assurance in the user's identity there is potential for a person to impersonate another or supply false contact information. This makes it difficult to always hold a real person accountable for the malicious actions on the system or theft of intellectual property.	M	L	5 Low	4
Cleartext authentication is used for mail list management.	Cracker	Majordomo email lists are managed by cleartext passwords in the emails and trust the source address for authentication. This opens up the possibility for eavesdropping, list denial of service, exposure of private archives, and list hijacking for spam.	M	M	25 Medium	5
Conference Call system uses weak or loosely managed PINs	Cracker	PINs or passwords aren't used for most conferences, and the call invitation is in the clear in emails and on attendee's calendars. This opens XSEDE up to the threat of eavesdropping on private meetings.	M	L	5 Low	6
Self-service password resets can be used by anyone able to intercept or monitor emails	Cracker	An attacker could try to reset a user's account from the portal, capturing the email with the reset code. Then they could log onto resources that the do not have access to normally and and/or act as	M	L	5 Low	7

		another person maliciously.				
Credential Management						
User credentials are not always encrypted on disk	Cracker or curious user	Kerberos tickets and proxy certificate keys are usually unencrypted on disk on the HPCs. File permissions and short life spans are the primary mitigation against credential theft in this case. Gateways such as GlobusOnline which store many more credentials have more protections in place.	L	L	1 Low	8
Some services use unencrypted private keys to access XSEDE resources	Crackers	Several services run on the backend transferring data to/from the XDCDB, IIS, RDR, IIS metrics, etc. These often depend on permanent SSH keys or GSI certificates to run in automated scripts. Depending upon the lifetime of these credentials and the restrictions on the corresponding accounts, an attacker who compromises one of these systems could steal credentials to leverage elsewhere in an attack on XSEDE infrastructure.	M	M	25 Medium	9
Users control their keys and may not protect them adequately.	Cracker	XSEDE doesn't control how users protect SSH private keys or keys corresponding to their X.509 certificates on their own systems. Keys may be unencrypted, or encrypted with poor passphrases. This means a compromise on a user's system could allow an attacker to steal their credentials to logon to an XSEDE resource and spread their attack.	M	L	5 Low	10
System Hardening and Resiliency						
XSEDE lacks centralized logging infrastructure	Cracker	There is no centralized logging for XSEDE services. This makes it easier for an attacker to erase his digital trail. It also makes it more difficult for incident response teams to investigate a complex attack.	L	L	1 Low	11

DNS system lacks authentication for response and synchronization.	Cracker	DNSSEC is not used for clients and inconsistently used for server synchronization. This lack of authentication could be exploited to give false responses or poison servers. Besides DoS attacks, this could allow man-in-the-middle attacks for non-certificate based services like regular SSHD and default Globus configurations that rely on DNS for hostname canonicalization.	M	M	25 Medium	12
XSEDE hardening guidelines for SPs are unaudited lending to the possibility weaker security at some SPs.	Cracker	XSEDE has developed new guidelines for system hardening, but these are not enforced nor audited for compliance. Because of the shared trust fabric, XSEDE is again only as strong as its weakest link, and an attacker could exploit a more lenient security posture at one SP to spread an attack on XSEDE.	L	M	5 Low	13
XSEDE has inconsistent or non-existent backup process for key resources.	Accident or incompetent employee	There are many services and systems distributed across XSEDE, but there is no centralized backup or backup policies for critical resources. The hosting usually SP determines what if anything is backed up. Equipment failure or a major security incident could make it hard to bring these systems back online in a secure state (especially since not all sites test their restore processes).	M	M	25 Medium	14
XSEDE does little to no centralized security monitoring.	Cracker	XSEDE networks and publicly facing systems do not utilize any sort of intrusion detection systems and could be compromised with delayed notice. Defacement or disruption of the portal, XSEDE's public face, would be potentially damaging, and lack of monitoring increases the exposure time during incidents. Given the large bandwidth and capabilities of HPCs and perfSonar nodes, a sizable amount of damage from a DoS attack using XSEDE resources could occur in even a	M	M	25 Medium	15

		short amount of time.				
Admins are not required to use strong authentication for management of XSEDE services.	N/A	RETIRED	N/A	N/A	N/A	16
A critical service is vulnerable to a denial of service attack.	Cracker	Any Internet facing system could be vulnerable to a DoS attack, given sufficient adversarial resources. Even though XSEDE tries to mitigate this by having replication across multiple sites and redundant network paths, it is still possible for an adversary to mount such an attack, especially against non-replicated services like the XUP and XRAS.	M	M	25 Medium	17
There is no consistent patch management process for all XSEDE services and systems.	Cracker	RETIRED. We are using Qualys and have a program in place now.	N/A	N/A	N/A	18
Privacy						
User data has weak isolation guarantees on most resources.	Curious Users	There is little besides file system permissions or ACLs that isolate users and their data on most XSEDE systems. There is potential for such basic mechanisms to be overcome and allow data snooping by adversaries. When networked filesystems are used without encryption, this threat is increased.	L	L	1 Low	19
Tickets are emailed in plaintext.	Cracker	Much of the ticket system communication is done over plaintext emails. Since some of the tickets are sensitive and contain security relevant information, attackers snooping those emails could gain an advantage.	M	L	5 Low	20
Incident response team	Cracker	The Jabber server uses SSL and does not log conversations.	L	L	1 Low	21

members may log sensitive IM chats.		However, there is no control over the endpoints when using the Jabber server for incident response. So messages could be logged on client hosts and accessed more easily or exposed if a laptop is lost. This inside information could be used for gain by an attacker.				
Software						
Deployed software could be out of date with stale CTSS registrations info.	Cracker	While there is a common software stack for XSEDE compatibility, there is an inconsistency across sites on versions deployed, with some sites using very out-of-date pieces software. This affects major software and central services less as we are using Qualys. However, there is a threat that an attacker could leverage an exploit at one site to gain a foothold at another.	L	M	5 Low	22
XSEDE relies heavily on in-house software that has not had code audits.	Cracker	Many applets and pieces of software for XSEDE have been developed in-house without code reviews or expertise in security. There are likely unknown security flaws that could be exploited in a targeted attack. This is especially true of something as complex as the user portal, whose compromise would be harmful to XSEDE's reputation.	M	M	25 Medium	23
Some XSEDE resources depend upon proprietary, unvetted protocols.	Crackers	Some proprietary protocols have been created for services (e.g., Globus listener over UDP). Developing secure protocols in nefariously hard, and in some cases there is no indication that any encryption or signing has been done. A very targeted attacker could exploit protocol vulnerabilities in ways that are very difficult to detect or deter.	L	L	1 Low	24
XSEDE depends upon software that is no longer actively	Cracker	XSEDE depends on some software that no longer has active development. This means there may be no one there to fix security or reliability bugs which could be	M	M	25 Medium	25

supported.		exploited maliciously (e.g., gaining shell access, privilege escalation or DoS), and it has caused migration challenges as we have had to move to new standards like SHA-2 certificates.				
Globus Online does not follow best practices with all of their key handling.	Cracker	RETIRED. Reasons: GO has significantly mitigated the risk of exploitation with recommended controls, and this is a more specific instance of #8	N/A	N/A	N/A	26
There is a zero-day root escalation exploit in the wild for Linux or some common piece of the XSEDE software stack	Cracker	Software vulnerabilities are commonly found, and there are are often crackers sitting on harvested user credentials waiting for the next Linux zero-day that would allow them to escalate their privileges to obtain root on an XSEDE resource.	H	M	50 Medium	27
There is a common XSEDE service with a remote exploit	Cracker	By virtue of having services online, there is always a risk that a new vulnerability is discovered that allows remote exploit that could either be combined with a local root escalation or that gives root itself. If such an exploit is in the wild and XSEDE is vulnerable, it is only a matter of time before scans by crackers discover it and exploit it.	M	M	25 Medium	28
Policy and Procedure						
Security policies and procedures from Teragrid are out-of-date.	N/A	RETIRED. XSEDE has updated policies and produced new ones.	N/A	N/A	N/A	29
Services and software grand-fathered into XSEDE have not gone through the formal review processes that are now a part of XSEDE	Cracker	With inconsistent standards, attackers could target grand-fathered services which might not pass security reviews today, though many of these are retired or replaced in time.	L	M	5 Low	30

There is no regular security training for XSEDE staff, and the annual training for new users is optional and only available at the XSEDE conference.	Mistakes and accidents	If users do not know what they should and should not do with respect to security, it is more likely that they can be taken advantage of by social engineers and that they do not protect their credentials well enough. Because we don't audit SPs, this makes it more likely that XSEDE staff are not following security baselines either. Failures to implement policy and procedure are therefore more likely and could result in more downtime and slower response to an incident.	H	M	50 Medium	31
Incident response resources at different SPs vary and the XSEDE IR team is geographically distributed.	Cracker	Different sites are more or less prepared to detect and respond to incidents, giving adversaries strategic options when deciding where to attack XSEDE. Being geographically and institutionally separated, the XSEDE IR team members must work harder to coordinate activities and overcome hurdles with information sharing that could slow response.	M	M	25 Medium	32
There are no security baselines for most services and there is no regular auditing with respect to security.	Cracker	As time goes by, systems likely drift from more secure and up-to-date configurations to less secure states. Crackers can target such infrastructure and find vulnerabilities to exploit more easily.	M	M	25 Medium	33
Systemic						
XSEDE has very complex organizational and procedural structures.	Accident or oversight	With a large complex organizational structure with many geographically separated employees only part-time committed to XSEDE, there is increased possibility for inaction or slow action at a critical moments. This is exacerbated when roles and process are unclear and has an impact beyond just security, though agility is particularly important for security.	M	M	25 Medium	34
Resources for	Mistakes or	RETIRED. Security operations has	N/A	N/A	N/A	35

security in XSEDE could be inadequate.	oversights'	remained will supported in the first 5 years and is a part of the renewal process				
New in 2015						
Passwords can be compromised and harvested	Crackers	Crackers can sit on passwords that are compromised and wait for privilege escalation vulnerabilities. These exposures may not even happen on XSEDE resources but bleed over to XSEDE because of password reuse.	H	M	50 Medium	36
Most SPs have not done risk assessments	Cracker	Because individual SPs have not completed similar risk assessments a level lower, XSEDE cannot fully understand its risks. We assume the physical and cyber security of these centers as a starting point, and we may miss important vulnerabilities that should be addressed	M	L	5 Low	37
XSEDE has no social engineering awareness training for staff.	Cracker	Social-engineering is addressed in the optional annual user training for users help at the XSEDE conference. However, staff aren't trained to protect against social engineering, and it isn't considered as a part of most processes. Mitigating this though, account resets are automated and don't require a person in-the-loop anyway, and resources are physically spread out across many organizations. Finally, open scientific research isn't often threatened by targeted attacks for data.	L	L	1 Low	38
Passwordless remote root keys used for XWFS	Cracker	The XWFS is configured in a way that requires remote root logins across sites via passwordless SSH keys for GPFS. While some mitigations have been in put in place, this could allow a root exploit to spread across sites undetectably.	L	M	5 Low	39

07: Control Analysis Matrix

In this table we list current or planned security controls with descriptions. The status column uses "E" for existing and "P" for planned. The Federation column uses "Y" for appropriate to call a control for the federation level risk assessment, "N" is for not appropriate, and "M" is for maybe. Some controls are in a gray zone. For example, they may need to be implemented at the SP level, but they should follow policy or baselines established at by the federation. For example, local passwords and public keys are used differently at different SPs, but that doesn't mean there shouldn't be any sort of baseline as this has an effect on the overall security of XSEDE beyond an individual site. The description column tries to point out these sorts of ambiguities.

A source for the list of these controls came from brainstorming, reference to the security controls catalog in the NIST Special Publication 800-53, and a review of previously identified controls in 04: XSEDE System Characterization.

Control	Status	Federation	Description
Access Control & Accounting			
Kerberos	E	Y	The XSEDE Kerberos infrastructure is mostly transparent to users, and many probably don't realize their portal password is really for a Kerberos principal. Kerberos is also used for the single-sign-on hub, the XSEDE MyProxy server, and O4MP service used by GlobusOnline. By using this control, account management can be centralized, and accounts disabled quickly, though existing certificates are not automatically revoked when a Kerberos principal is disabled.
PKI	E	Y	The XSEDE PKI includes the IGTF accredited MyProxy CAs used by the OA4MP service and SSO Hub, the CI Logon infrastructure, the trust store of approved XSEDE root certificates, all the existing short-lived certificates for XSEDE users, and the set of policies and procedures for vetting CAs for the tarball. The policies for how the myproxy.xsede.org and CI Logon CAs are managed follow from the CPS approved by TAGPMA. The purpose of the XSEDE PKI is to allow single-sign-on and use of federated identities and to support various grid middleware services. It is another type of authentication control that lets XSEDE control who can authoritatively vouch for user identities. XSEDE no longer manages a PKI for host certificates.
Certificate settings	E	Y	Certificate lifetimes and restricted proxy settings are important parameters to tune risks with regard to the XSEDE PKI. Certificate lifetimes are limited by the CA policies and often even shorter by configurations for services that request certificates with shorter lifetimes. Restricted certificates for GridFTP only are used by GlobusOnline. While this is all a part of the PKI, it deserves special attention as a set of "knobs" on that control.

OAuth Service	E	Y	Rather than having external portals request users enter XSEDE Kerberos password in their web sites, XSEDE provides an OAuth front-end to sites like GlobusOnline. This control protects against exposure of credentials at non-XSEDE sites and trains users to only enter XSEDE credentials at the XSEDE portal. This control centralizes web-based authentications for XSEDE users.
SP passwords, OTP & Keys	E	M	Users often don't login directly to XSEDE compute resources with their XSEDE Kerberos principal, except through the SSO Hub. Each site handles authentication differently with different technologies and policies. Because of the single-sign-on capabilities provided by linking these authentication mechanisms with local CAs, one could argue XSEDE authentication is only as strong as the weakest link. Therefore, these local authentication systems are relevant to XSEDE (1) when they can be used to get a certificate from a trusted CA and (2) when they are used for administrative interfaces of shared critical infrastructure. Standard have been made that require two-factor for administrators of XSEDE resources, but there are no official standards on passwords or use of keys.
File System Permissions & ACLs	E	M	Protection of private keys and user files is done via file system permissions and ACLs. Configurations are unique to different hosts administered by different people at institutions. For global filesystems and shared critical support services, this is definitely a federation-level access control. It is less of a federation level-issue for individual SP compute and storage resources, except regarding how certificates are protected. The latter should be standardized with a new administrative control policy.
Self-service Password resets	E	Y	While mostly a usability feature, self-service password resets can encourage users to choose better or unique passwords. Security benefits really depend on how they are configured and the alternatives the are compared against. It is debatable whether or not this is a security control.
Policy & Procedure			
Security Policies	E	Y	There are several security relevant policies, standards and procedures at www.xsede.org/security/.
User Provisioning & Vetting	E	Y	Specific procedures (which were approved by TAGPMA) have been produced for vetting users and allocating accounts. These processes are important for both security and accounting. Changes to them should be carefully considered as they form the founding assumptions for much of the trust fabric within XSEDE, and the effectiveness of other controls depend upon them.

User Training	E/P	Y	Security operations has developed some using training materials for security. However, there is no mandatory or targeted security training in XSEDE. There are plans to develop more security training materials in the future and to find new methods of outreach. The effectiveness of this control largely depends upon the ability to reach the people who most affect the security of XSEDE. This may be developers, system administrators and other XSEDE staff more than users.
Change Control & Testing Process	E	Y	The Systems, Software & Engineering group oversees the process of making configuration changes to the baselines held by Operations. There are official processes for making CI changes consisting of reviews in SD&I and the Software Testing & Deployment (STD) group. All changes are supposed to go through reviews and testing. This is a more general quality control with security implications.
Security Reviews	E	Y	Configuration Items (CIs) are supposed to go through security reviews in both SD&I and Operations before they are deployed by the STD group. The purpose of this control is to preemptively detect weak configurations or designs.
Risk Management Program	E	Y	XSEDE has a formal, risk-based security program. This risk assessment is part of a larger risk-based program designed to more judiciously apply resources to security. This is an administrative control.
Incident Response Program	E	Y	XSEDE funds professional incident responders at multiple institutions who can quickly investigate security problems and work with their contacts at the service providers. This control is focused on containment and resolution after a failure of other controls.
Security-Related Activity Planning	E	Y	XSEDE has formal security related planning activities and resources to deploy new security controls.
Information System Component Inventory	E	Y	XSEDE maintains a master registry listing all the services provided, ranking them into tiers based upon criticality and service-level promises. This is an important list to have and keep updated as it tells us where to focus monitoring and vulnerability management efforts.
Security Production Baseline arch for AuthN/Z	E	Y	XSEDE does have a formal architecture and baseline for one set of security controls, namely those core technologies related to authentication and authorization. This architecture and baseline is important to refer to when integrating any new technologies so that we can keep a consistent and unified authN/Z infrastructure. This control is necessary with the complexity and sheer number or organizations involved.
Intelligence Channels	E	Y	In order to share information between sites and so that all members of the security working group could subscribe to intelligence from REN-ISAC feeds, it is a requirement for member organizations to join REN-ISAC.

Resiliency			
Offsite Redundancy	E	Y	Most all of the systems and support services from the master registry are replicated across at least two service providers. This protects against hardware, environmental and some organizational failures.
Backups	E	M	Each site is doing backups in its own way or not at all. There is certainly an XSEDE stake in doing backups for shared services. Though there is not any current guideline or policy on that, replication mitigates the need to an extent.
Other Logical Controls			
Transport Layer Encryption	E	Y	Most services encrypt data in transit and by using SSH, HTTPS or other proprietary protocols over SSL. This is strongly encouraged or required in the security reviews of new CIs especially to protect user credentials, but also to authenticate services to users and each other.
Previous Login Notification	E	M	Most systems tell users when and from where they previously logged in. This control helps detect compromised accounts.
Firewalls	E	M	XSEDENet does not use firewalls. There are no centrally planned host or network-based firewalls for XSEDE services. Some of the SPs have a combination of both host and network firewall services, but these are not documented. XSEDE has guidelines or hardening requirements for XSEDE central services that include rules about host-based firewalls.
Audit Logging	E	M	There is no centralized syslog server or log management system. Each service and each SP handles this differently. XSEDE has a baseline standard for audit logging, but this just sets the requirements and not the details of implementation.
Constrained Database views	E	Y	There are several support services that utilize databases (e.g., AMIE). Some of these provide separate database views, employing the principles of least privilege to expose only the necessary pieces with the necessary rights to particular systems and users. This is a practice encouraged in design reviews of new services.
Secure Communicatio n for IR	E	Y	The incident response team is distributed and therefore relies upon secure communication channels. They have their own private teleconference code, secure IM services, and a secure emailing service with PGP.
Secure IR Collaboration Space	E	Y	The incident response team has its own restricted wiki to share information about current incidents.

Vulnerability Scanning	E	Y	Security operations has a vulnerability management program for XSEDE resources based on QualysGuard. This control is complimented with an administrative control in the form of policies about how quickly different levels of vulnerabilities must be remedied or exceptions must be submitted.
Trusted Network Path	E	Y	XSEDENet is a private network backbone just for communication between major XSEDE SPs. Regular traffic is not routed to transit over it. Some sites trust this more and apply looser rules for firewalls and monitoring.
IDS	E	N	Most level 1 SPs use Bro for an IDS, and there is strong Bro expertise and support in the community
Virtualization	E	Y	XSEDE utilizes virtualization for many centralized support services, many of which are hosted at IU. While not done directly for security, this allows strong isolation of services, quicker recover, and the ability to replicate them more quickly if there is a DoS or other attack.
Spam Filters	E	Y	XSEDE has several mail lists all of which have some spam filtering done on the server side.

08: Control Recommendations

XSEDE should begin with implementing the controls labeled high priority that address the three risks ranked 50. Next they should move onto the medium priority risks. Prioritization within that class should be based on how many significant risks a control addresses, how effective the control is expected to be, and how much the control costs in terms of time, resource or political will.

High Priority (Risk(s) Addressed >= 50)

1. RECOMMENDED Change password reset procedures to choose random passwords for users.
    - Mitigates risk #27 and #36
2. RECOMMENDED Develop and rollout mandatory security training for XSEDE staff. The goal is to familiarize them with important policy and procedure and make them resistant to social engineering.
    - Addresses risk #38 and mitigates #13, #21, and #31.
3. Require two-factor for all XSEDE users
    - Addresses #27 and #36
4. Provide simple, light-weight user profiling tools for SPs to detect compromised accounts before they are used in conjunction with a zero-day exploit.
    - Mitigates risk #27, #36, #4, #7, #10

Medium Priority (Risk(s) Addressed 25-49)

1. RECOMMENDED Maintain strong leadership and regular meetings for XSEDE security ops. Participation of level 1 SPs should not be optional. Security personnel and resource allocations should also be evaluated and redistributed at least annually to be proportional to where the work is done or needed to be done.
    - Mitigates risk #27, #28, #32 and #34
2. RECOMMENDED Perform audits to make sure XSEDE systems and services stay true to baselines. Audits need not be invasive and may consist in part of questionnaires to be filled out by SPs or checklists for them to complete.
    - Addresses risk #13 and Mitigates #11, #14, #22, #31, #33, and #37
3. RECOMMENDED Have regular security incident drills and require participation from all level 1 SPs.
    - Mitigates risk #32 and possibly #38
4. RECOMMENDED Develop per service security baselines for critical XSEDE support services.
    - Mitigates risk #13 and #33
5. RECOMMENDED Identify unsupported software in the XSEDE software stack and either (a) retire and eradicate it, (b) update or replace it, or (c) commit resources to adopt and maintain it.

    - Addresses risk #25

6. RECOMMENDED Identify services not replicated across at least two XSEDE SPs and either replicate them or move to a service provider with DoS protections. Focus on critical services first.
    - Mitigates risk #17
7. RECOMMENDED Identify and minimize services running as root on XSEDE systems; reconfigure to run as non-root wherever possible and recommend jails or other containment mechanisms where this is not possible.
    - Mitigates risk #28 and potentially #39
8. RECOMMENDED Inventory how system accounts are used for various services and identify those using unencrypted certificates or SSH keys. Protect the credential as best as possible, limit the capabilities of the credential and corresponding account to only those needed, and monitor for any unexpected use of such accounts.
    - Mitigates risk #9 and maybe #39
9. RECOMMENDED Either internally or using an external service, perform code audits of commonly deployed XSEDE software. This could mean requiring them to use SWAMP or coverity as part of the SD&I review.
    - Addresses risk #23 and mitigates risk #30.
10. RECOMMENDED Deploy DNSSEC across the xsede.org domain.
    - Addresses risk #12.
11. Utilize a redundant, offsite backup service for critical XSEDE services. This is not for user data, but system resiliency.
    - Addresses risk #14.
12. Deploy a central syslog server and host-based IDS for XSEDE as well as correlate network IDS data from sensors on XSEDENet.
    - Addresses risk #15 and #11
13. Identify grand-fathered in services and run through SD&I review process.
    - Addresses risk #24 and #30

Low Priority (Risk(s) Addressed < 25)

- RECOMMENDED Migrate from majordomo to an email list service that requires login over HTTPS for list and subscription management, rather than plaintext emails with plaintext passwords for management.
    - Addresses risk #5
- RECOMMENDED Create a standard for authentication to receive an XSEDE acceptable certificate and audit for compliance.
    - Mitigates risk #2

- RECOMMENDED Identify and remove local or non-XSEDE authentication systems for XSEDE services and replace with XSEDE Kerberos or GSI where possible.
    - Addresses risk #3
- RECOMMENDED Require a second email address to be registered for account creation.
    - Mitigates risk #7
- Rework the processes of account creation for XSEDE to reach a higher level of assurance, likely relying on email and the PI less for identity vetting.
    - Mitigates risk #4
- Require authentication for Lync meetings.
    - Addresses risk #6
- Utilize security questions for self-service resets falling back to manual processes if that fails.
    - Addresses risk #7
- Shorten ticket or certificate lifetime
    - Mitigates risk #8.
- Disable publickey SSH access on XSEDE resources and manage certificates for users securely on their behalf.
    - Addresses risk #10
- Use network filesystems that encrypt in transit for user filesystems.
    - Mitigates risk #19
- Use encrypted filesystems with strong guarantees of isolation between users.
    - Addresses risk #19.
- Do not put ticket contents in the emails but only a URL where one must authenticate to add to the ticket or view it.
    - Addresses risk #20
- Require full disk encryption for incident response team members workstations and laptops.
    - Mitigates risk #21

Appendix A. XSEDE Systems & Services

This appendix contains details about the various services shown in the Systems & Services table of section 04. XSEDE System Characterization. This list was accurate for the 2012 XSEDE Risk Assessment.

A.1 Level 1 SP Compute and Storage Resources

The most visible and arguably most critical XSEDE resources are the major compute and storage resources provided by the level 1 service providers (SPs). These are all tightly integrated with common authentication systems and the high-speed XSEDEnet backplane. While XSEDE would not fail with any one of these resources absent, there would be little to XSEDE if they were all unavailable to our user community.

COMPUTE RESOURCE NAME	HOST SP
Blacklight	PSC
Gordon CC	SDSC
Gordon ION	SDSC
Trestles	SDSC
Stampede	TACC
Comet	SDSC
SuperMIC	LSU
Darter	NICS
Mason	Indiana
Wrangler	TACC
Maverick	TACC
Nautilus	NICS
OSG Condor Pool	USC

STORAGE RESOURCE NAME	HOST SP
Ranch	TACC
HPSS	NICS
Wrangler	TACC
SuperCell	PSC
Data Oasis	SDSC
XWFS	TACC/NICS/NCSA

The most current list is at https://www.xsede.org/resources/overview .

Data

There should be no PII (Personally Identifiable Information) or other legally protected data associated with XSEDE allocations on these systems. Nor should there be proprietary company data under NDA or other restrictions as the XSEDE allocations are for open scientific research funded by the NSF. However, this

does not mean that individual SPs do not share their resources with private sector partners who may have proprietary data on the systems. However, that is not an XSEDE issue, and it would be between the individual SP and their customers how they protect such sensitive data.

This said, the data is still not public, and XSEDE provides confidentiality to users through use of file-system permissions and ACLs. Furthermore, a given compute node is usually only running one user's job at a particular moment, and other users cannot typically login to another's active compute nodes. While XSEDE SPs strive to deliver reliable storage providing for integrity and availability of user data, the ultimate onus is on individual users to make sure they backup their data.

In addition to user data for compute jobs, there are credentials stored on these systems. Kerberos tickets or GSI certificates allow for single-sign-on, but this means that private keys are sometimes stored on disk. In such cases they are protected by restricted file system permissions. Users may also create SSH private keys and put them on systems manually, but it is up to them to protect these keys they generate themselves. Some SPs don't allow public key authentication, and others check that users encrypt the keys with passphrases.

Data Flow

Primarily, user data is moved manually gridftp or sftp protocols (often through a front-end like GlobusOnline). These services typically use either GSI certificates or Kerberos for authentication, but some allow public key authentication with SFTP.

Most SPs mount distributed filesystems that abstract data movement over the network from users. For example, the XSEDE-wide file system is mounted on most of these compute resources.

Credentials typically don't flow from system to system, but instead credentials on a client system are used to get new proxy credentials on the remote system. Gridmap files and other information related to authorization and accounting are pushed from the XSEDE central database are discussed elsewhere.

System Interfaces

Gridftp, SFTP and XWFS are discussed in the section above on data flow. In addition to these interfaces for data flow, there are interfaces for shell login (i.e., ssh, gsi-ssh), job submission (e.g., GRAM and Unicore), accounting (e.g., AMIE which is discussed later), and more. These and others are discussed more thoroughly in the section on the XSEDE software stack.

In addition to the core services listed above, there is nothing preventing individual SPs from running various other services, but these are not XSEDE interfaces and are more appropriately documented in risk assessments at individual SPs.

A.2 AMIE

More information is available on AMIE at http://software.xsede.org/production/xdcdb/amie-docs.tar. Documents found there include AMIE Model, AMIE Installation and Configuration and Implementing AMIE documents. The AMIE code itself is at http://software.xsede.org/production/xdcdb/amie.tar.

The AMIE model consists of two sites and an agreed upon set of transactions that the two sites will use to send account management data to each other. A transaction consists of packets of data sent between the two sites. The site sending a packet is called the local site and the packet is known to the sending site as an outgoing packet. The receiving site is called the remote site and the packet received is known to the receiving site as an incoming packet. The site that creates the transaction (and sends the first packet) is also called the originating site.

AMIE is transaction based. Transactions have a number of properties. These are the local site, the remote site, the originating site, a transaction id, and a state. Once created, the first four properties do not change. However, the state changes over time.

The transaction id is used to distinguish one transaction from another. The originating site chooses the transaction id without consulting the remote site. The only rule is that a transaction id created by one site may not be reused by that site for a different transaction. Hence a transaction is identified by the originating site, the transaction id, the local site, and the remote site.

AMIE defines 3 states for a transaction: in-progress, completed, or failed. The initial state of a transaction is in-progress. It remains in that state until all packets have been processed. If all packets have been successfully processed, the transaction state becomes completed. If any of the packets causes a failure, the transaction state becomes failed.

Transactions also have packets which contain account management data. Incoming packets are those packets received from the remote site. Outgoing packets are those packets created by the local site to be sent to the remote site. Outgoing packets are created either when the transaction is created or as a reply to an incoming packet.

AMIE does not specify a pre-defined set of transactions. It specifies a set of packets which can be used within transactions, but the sites must agree on the packets that are used within transactions as well as the ordering of those packets.

A packet has a number of properties. These are type, version, packet id, and state. It also has a list of expected replies.

The types and their contents are determined by the AMIE XML schema. As of the writing of this document, the AMIE XML schema only specifies version "1.0". However, in anticipation of newer versions, each packet must also specify its AMIE version, since the content of the packet depends on the version.

Each packet has a packet id which is chosen by the site that creates the packet. It has to be unique within the set of outgoing packets for a given transaction.

Data

AMIE accounting transaction data relies on information in the XDCDB and information sent back and forth to XSEDE service providers. Integrity is much more important than confidentiality of this accounting data. Still, much of it is not public information

Data Flow

AMIE has a centralized service that all Service Providers communicate with, and AMIE has services run at each SP. Data flows back and forth between the XDCDB and databases at the local sites through AMIE transactions.

System Interfaces

AMIE is a transaction based system with interfaces (tunneled over SSH) at each Service Provider and an XSEDE AMIE interface in front of the XDCDB.

A.3 CILogon

The CILogon Service provides secure access to XSEDE services using InCommon campus logins. It obtains short-lived credentials from a private CILogon MyProxy server for user sessions. The CILogon Service provides an alternate mechanism of authenticating to Globus and XSEDE services and does not replace existing authentication mechanisms.

See http://www.cilogon.org/xsede and http://www.cilogon.org/portal-delegation for details.

Data

The CILogon service maintains a database of user authentications to track name/email changes which would affect the certificate subject Distinguished Name (DN). No passwords are stored by the CILogon Service.

Data Flow

XSEDE services request user data from the CILogon Service using the OAuth protocol. Users authenticate with their campus identity providers, which release SAML attributes to the CILogon Service. The CILogon Service then issues short-lived credentials to XSEDE services for use on behalf of the users.

System Interfaces

User and service access is via HTTPS (port 443). SSH access for system administration is restricted to bastion hosts secured with two-factor authentication.

A.4 CI Tutor

www.citutor.org at NCSA. Contact: Sandie Kappes <kappes@illinois.edu>. CI-Tutor provides the means to learn about High Performance Computing and CyberInfrastructure (CI).

Data

CI-Tutor can be accessed by anyone, and thus there is no sensitive information. That said, users still need to create an account with a username and password. Passwords are stored in an encrypted form in a MySQL database.

Data Flow

CI-Tutor is not currently hosted on an NCSA or XSEDE server, but is hosted by SiteGround which is a cloud based hosting service. It is a completely isolated platform with its own authentication system, and no data flows between it and other XSEDE resources.

System Interfaces

HTTPS is the only interface provided to XSEDE, though SiteGround could of course have its own special administrative interfaces abstracted away from us by the cloud service.

A.5 Virtual Workshop

https://www.cac.cornell.edu at Cornell. (Linked to from XUP) Contact: Resa Alvord <rda1@cornell.edu> .

Data

Nothing sensitive to XSEDE is stored here. It is all educational materials, available to anyone who registers an account. This is done automatically for XSEDE users who follow from the link on the XSEDE portal.

Data Flow

It is it's own isolated system, but XSEDE users are pre-registered if they come from the portal. This is simply the XSEDE username being passed as an HTTPS POST.

System Interfaces

XSEDE users only interact with the service over HTTPS. Cornell has it's own administrative interfaces most likely, but being an isolated resource provided by a third party, there is little concern about XSEDE being affected as no XSEDE credentials are stored on this service.

A.6 XSEDE.org DNS

HOSTNAME	LOCATION	ADMIN	COMMENTS
ns1.xsede.org	NCSA	netneng@ncsa.illinois.edu	
dns1.tacc.utexas.edu	TACC		
dns2.tacc.utexas.edu	TACC		
dns3.tacc.utexas.edu	TACC		

XsedeNet has multiple DNS servers currently in production. The primary DNS (ns1.xsede.org, 141.142.143.137) server is located at NCSA. In addition, there are three DNS servers that serve as secondary DNS located at TACC. Also, there are several secondary DNS delegations throughout the XSEDE network.

XSEDE Secondary DNS Delegations

ZONE	NAME SERVER	CONTACT	DELEGATED
NICS	ns0.nics.utk.edu, ns1.nics.utk.edu	Stephen McNally ,<smcnally@utk.edu>	yes
IU	dns1.iu.edu, dns2.iu.edu	dns-admin@indiana.edu	yes
SDSC	ns0.sdsc.edu, ns1.sdsc.edu	noc@sdsc.edu, hutton@ucsd.edu, jeff@sdsc.edu	yes
NCSA	dns1.ncsa.illinois.edu, dns2.ncsa.illinois.edu	neteng@ncsa.illinois.edu	yes
PSC	dns1.psc.edu, charon.psc.edu	pscnet-admin@psc.edu	yes
TACC	(dns1,dns2,dns3).tacc.utexas.edu	jones@tacc.utexas.edu	yes

Data

All data is public and consists of just CNAME and A records.

Data Flow

Data is replicated over DNS protocol tcp/53.

System Interfaces

Primarily the DNS protocol. SSH access is used by admins to update configuration and tables manually.

A.7 Conference Call System

XSEDE utilizes the University of Illinois's Unified Communications Lync service for video and audio conferencing.

Data

Discussions can be very sensitive, but they are not archived.

Data Flow

The URLs can be sensitive since often they are all that is needed to join a call. People generally email these and/or put them in meeting invitations.

Credentials are not needed, though UIUC people tend to login with their UIUC netid. Few if any calls use PINs or extra protections, but those that do have different mechanisms of dissemination

System Interfaces

The conference call system works over traditional phone lines, a web client or Lync.

A.8 Mail Lists

HOSTNAME	LOCATION	ADMIN	COMMENTS
relay-*.ncsa.uiuc.edu	NCSA	Chris Lindsey cpl@illinois.edu	
pop.ncsa.uiuc.edu	NCSA	cpl@illinois.edu	
zimbra.xsede.org	NCSA	cpl@illinois.edu	
mhonarc.xsede.org	NCSA	cpl@illinois.edu	

The mail lists (both private and public) for XSEDE are run by the NCSA majordomo mail list server relays. A hot backup sits on the VM farm at IU, and a Zimbra server is used to allow applications to retrieve mail, share files securely, and provide IMAPS/POPS access for applications.

Data

Besides membership data in majordomo config files, the archived lists create data. All email lists are archived, and no distinction is made based on sensitivity.

Some files are shared through the Zimbra briefcase feature, though nothing containing PII or specially regulated data.

Data Flow

Most data flows as email through over SMTP. Searchable archives are accessed over HTTPS. Email read by applications is done over SSL, and the Zimbra syncs files over SSL.

Email lists are managed by plaintext passwords sent in email. So these credentials are sent around unencrypted, which if discovered, could add someone to any list and hence the archives

System Interfaces

Most management and use of the lists is by email over SMTP. Archives can be accessed over HTTPS. System admins configure majordomo through the command line over SSH. Applications access Zimbra through POPS or IMAPS.

A.9 Genesis II

HOSTNAME	LOCATION	ADMIN	COMMENTS
gffs-{1,2}.xsede.org	IU	Mike Lowe (jomlowe@iu.edu)	
sts-{1,2}.xsede.org	IU	Mike Lowe (jomlowe@iu.edu)	

The Genesis II rootname service (RNS) and secure token service (STS) are deployed in support of GFFS on XSEDE resources. The STS plays an analogous role to MyProxy and Oauth with the Globus Toolkit managing credentials and security tokens. The RNS service provides directory services, which could be used for GFFS (e.g., similar to AFS's global name space).

Data

- The GFFS protects data against unauthorized access using a system of ACLs (Access Control List) that control accessibility of every resource in RNS space.
- Data in transit is encrypted by the TLS protocol.
- Server-side data is protected by existing Operating System mechanisms, in that the container database is in a user's private folder. Passwords are not stored in clear text in the database.
- No guarantees can be made about data privacy once it reaches job processing (EMS), since those systems are outside of GFFS control.

Data Flow

- Users are authenticated against STS (Secure Token Servers) established in the grid.
- Client to server communication is encrypted with TLS.
- Users can authenticate against XSEDE MyProxy server to obtain their session identity for TLS. This login is also vetted against the XSEDE Kerberos server via a Kerberos STS in the grid.
- Trust delegation in GFFS relies on SAML trust delegations embedded in the SOAP headers, as defined in SDIACT-110 (https://jira.xsede.org/browse/SDIACT-110)
- User credentials are stored in local user state directory (by default in “$HOME/.genesisII-2.0”).
- Server-side records for STS identities (and all other state items for the GFFS) are held in the container's private state database.
- GFFS supports authentication via MyProxy, InCommon ECP, Username/Password, Kerberos, and X509 Keystore.

System Interfaces

- The GFFS provides a web services interface over TLS protocol, by default on port 18443.
- The GFFS Container is the software component that provides web services. It runs as a non-privileged user.
- To use the resources (GFFS and EMS), user has to have xsede account that can be authenticated with xsede myproxy and kerberos authentication.

A.10 Globus Listener

http://globus-usage.xsede.org at ANL. Contact: Stu Martin <smartin@mcs.anl.gov>. The Globus listener collects metrics on GTK services from the different SPs.

Data

Nothing is very sensitive. It just collects statistics on file sizes and endpoints used in different gridftp transfers.

Data Flow

GRAM/GridFTP servers send data to the globus listener service running at the IU VM farm.

System Interfaces

Globus listener uses a proprietary application layer protocol over UDP.

A.11 Globus Online

https://globusonline.org at ANL. Contact: Steve Tuecke <tuecke@ci.uchicago.edu>. Third party front-end for data movement through gridftp protocol between endpoints both within XSEDE and external to it.

Data

The service is used to transfer data between systems, most of which is neither public nor highly sensitive in XSEDE. It is primarily used for transferring large scientific data sets. Globus Online does hold short-lived end-user certificates on behalf of the users to make transfers. These certificates are just for GridFTP and do not have the rights to login through GSISSH.

Data Flow

Bulk data transfers are between GO endpoints which are either gridftp servers registered with Globus Online or desktop clients running GlobusConnect software.

GO manages the certificates and private keys for users necessary to make transfers. It receives the keys and certificates from the XSEDE MyProxy CA by passing an OAuth token. It gets the OAuth token by redirecting the users OAuth portlet on the XSEDE user portal

System Interfaces

GlobusOnline has a web interface used for most functions. There is also a CLI interface for scripting, and this requires SSH keys to registered on the GO website. Finally, the GlobusConnect client can be run on desktop systems to create another endpoint for transfer on your local machine.

A.12 Metrics for IIS

info-metrics.teragrid.org at IU. Contact: JP Navarro <navarro@mcs.anl.gov>. The IIS Metrics server collects all the logs for the info.teragrid.org and repo.teragrid.org/software.teragrid.org/software.xsede.org servers and processes those logs to produce metrics information.

Data

The logs are not particularly sensitive, just Apache and Tomcat logs.

Data Flow

Log data from info.teragrid.org and software repositories (repo.teragrid.org/software.teragrid.org/software.xsede.org) flow to the IIS metrics server. This is done with scp and public key authentication.

System Interfaces

SSH is the only interface for all transactions.

A.13 Integrated Information Systems

HOSTNAME	LOCATION	ADMIN	COMMENTS
info.xsede.org	NA	JP Navarro <navarro@mcs.anl.gov>	Rotating between info1 & info2
info.dyn.xsede.org	NA	JP Navarro <navarro@mcs.anl.gov>	Rotating between info1 & info2
info.dyn.teragrid.org	NA	JP Navarro <navarro@mcs.anl.gov>	Rotating between info1 & info2
info1.dyn.teragrid.org	IU	JP Navarro <navarro@mcs.anl.gov>	
info2.dyn.teragrid.org	Serveraxis.com	JP Navarro <navarro@mcs.anl.gov>	
info1.dyn.xsede.org	IU	JP Navarro <navarro@mcs.anl.gov>	
info2.dyn.xsede.org	NICS	JP Navarro <navarro@mcs.anl.gov>	

XSEDE information services are an integrated collection of web services (in the broadest sense) that publish information about the capabilities (systems, software, and services) available to the XSEDE user community. XSEDE Information Services defines the standards and services that that enable XSEDE resource and service providers to publish and advertise their offerings, via software interfaces, to the user community. The XSEDE user portal, user documentation, science gateways, and user applications can query information services software interfaces to discover information about XSEDE.

Data

IIS aggregates and publishes these types of information:

- Service provider scheduler jobs, percent loaded, and related batch system information
- Resource characteristics
- Resource capabilities, software, and service information
- Resource outages
- Resource GridFTP service speed information
- Capability descriptions
- Science gateway descriptions
- Resources accessible to each allocated project mapping

Most information is public, but some job information requires authenticated users. Nothing would be considered highly sensitive or legally protected PII.

Data Flow

Information flows FROM these sources to IIS:

- XSEDE central database, Speedpage database
- Resource Description Repository service
- Service provider information services
- NCSA HPC Software Catalog

Information flows FROM IIS to:

- XSEDE user portal and documentation
- Science Gateways
- INCA monitoring system
- Users via xdinfo command line tool
- Custom user or service provider applications

System Interfaces

It is all web interfaces (HTTP/HTTPS) whether through the xdinfo client or a browser. Admins also access it through a web interface.

A.14 INCA

Inca.xsede.org at SDSC. Contact: <inca@sdsc.edu>. Inca is a user-level, centralized monitoring system that runs periodic tests to determine the state of XSEDE provided services and software, in particular, services that level 1 & 2 SPs have registered with XSEDE as being provided. Inca also monitors GRAM usage, CA and CRL validity, and resource registration in MDS. Communication between the central server and the XSEDE resources occur using standard GSI credentials and SSL. The results of the tests are stored on the virtual machine inca.xsede.org. No information stored is particularly sensitive.

Data

Inca generates results for the following XSEDE resources (as of February 2015):

SITE	MACHINE
IU	Mason
IU	Quarry
IU	Software
LSU	Supermic
NICS	Bobo
NICS	Darter
NICS	Nautilus
OSG	Grid
PSC	Blacklight
SDSC	Capac
SDSC	Gordon
SDSC	Trestles
TACC	Maverick

TACC	Stampede

Data Flow

Monitoring information is collected from a regular user account called 'inca'. Our Inca server runs on a OTP protected VM machine at SDSC called capac.sdsc.edu. It spawns client daemons via SSH on the login nodes using our 'inca' account, which connect back to the server via a SSL connection. The client daemon, called the reporter manager, runs a number of user-level tests on a regular schedule via a perl cron module. When a test completes, the results are securely transmitted back to the central server. When it executes a test that requires a x509 proxy, it contacts the Inca server via the SSL connection and requests MyProxy information which it uses to download a short-term proxy credential. When it finishes running the test, it deletes the user proxy (http://inca.sdsc.edu/releases/2.6/guide/userguide-incat.html#PROXIES).

System Interfaces

GSISSH is used to run remote jobs and collect data. A web interface is used to view the data. The data generated by the Inca tests are world viewable at inca.xsede.org.

A.15 Jira

Jira.xsede.org at SDSC. Contact: Shava Smallen <ssmallen@sdsc.edu>. The XSEDE Jira system is used for task tracking for multiple XSEDE project areas, primarily SDI activities, though. It resides on the software.xsede.org host at the IU web farm.

Data

The ticket system data is stored in a database on the same system as the web server. There is little sensitive information in these tickets and little in regards to PII except the full name of users.

Data Flow

Tickets can be entered directly into the ticket system via the web interface, and comments can be added via email or the web interface. Actions on tickets may also generate emails back to the original reporter.

System Interfaces

All Jira interfaces are over HTTPS or email. The OS management interface is the same as that for software.xsede.org, described elsewhere.

A.16 Karnak Predictor

http://karnak.xsede.org at IU. Contact: Warren Smith <wsmith@tacc.utexas.edu>. Service provides predictions and information about jobs on batch scheduled computer systems.

Data

The backend database contains information about jobs and the current and historical ordering of jobs in batch queues. The job information includes usernames, but otherwise nothing particularly sensitive.

Data Flow

The service can be accessed directly by users as well as by other tools and services. The Karnak service is most commonly accessed via the XSEDE portal through Kerberos authentication.

System Interfaces

The Karnak service is a REST service at http://karnak.xsede.org. It currently provides content in HTML, XML, and text.

Users do not authenticate to the Karnak service and the services are not accessible over HTTPS. The service does not disclose user names or detailed job information so there is no need to secure client connections. Karnak receives information about jobs from the TeraGrid Integrated Information Services. It retrieves this by authenticating to (via host X.509 credentials) and then querying the WS-MDS service on info.teragrid.org.

A.17 XSEDE Kerberos Realm

The XSEDE Kerberos service provides username/password authentication for all XSEDE Portal users. It is a critical resource as the portal requires it as well as anyone who needs to use myproxy.teragrid.org.

HOSTNAME	LOCATION	ADMIN	COMMENTS
kerberos.teragrid.org	NCSA	Chris Lindsey @ NCSA	
kerberos-1.teragrid.org	NCSA		
kerberos-2.teragrid.org	PSC	Shane Filus @ PSC	

Data

The Kerberos servers contain usernames and hashed passwords for all XSEDE users.

Data Flow

The Kerberos service is primarily used by the XSEDE User Portal, myproxy.teragrid.org, and myproxy.psc.teragrid.org.

System Interfaces

Kerberos protocol and SSH through two-factor bastion for OS administration.

A.18 Certificate Authorities

HOSTNAME	LOCATION	ADMIN	COMMENTS
myproxy.teragrid.org	NCSA	ca-admin@ncsa.illinois.edu	
myproxy.psc.teragrid.org	PSC	ca-admin@psc.edu	

XSEDE accepts user certificates from the CAs listed at https://www.xsede.org/security. New CAs must be accredited by the International Grid Trust Federation. The XSEDE MyProxy CAs issue short-lived user certificates. These user certificates allow XSEDE users to authenticate to XSEDE grid services (gsissh, gram, gridftp). Some other CAs certificates are valid for up to 13 months, though not those run by XSEDE. Host certificates are provided by the InCommon IGTF-accredited CA, run by Comodo.

The myproxy.xsede.org server (primary, at NCSA) and myproxy.psc.xsede.org server (backup, at PSC) are used by the XSEDE User Portal and other XSEDE services to obtain short-lived certificates for XSEDE users who authenticate with their TERAGRID.ORG Kerberos username/password (i.e., their XSEDE Portal username/password). These servers use the NCSA MyProxy CA and PSC MyProxy CA on

the back-end. These CAs are critical to the full functionality of the user portal and GlobusOnline for data transfer.

Data

Each CA has a private key that it uses to sign certificates. The private key must be protected from disclosure. Compromise of the CA private key allows the attacker to create arbitrary certificates and thereby impersonate users and hosts.

Each CA gathers information to identify users and hosts for issuing certificates. Personally identifying information must be protected from disclosure.

Each user and host has a private key corresponding to their certificate that must be protected from disclosure. This is typically done with filesystem permissions.

Data Flow

Subscribers submit certificate requests to CAs which return signed certificates. These are public documents, protected from tampering by cryptographic signature. They do not require further protection.

When authenticating to MyProxy CAs, users provide their password to the MyProxy server for verification. Passwords are not stored on the MyProxy server. Connections are protected with SSL.

System Interfaces

MyProxy CAs use the MyProxy Protocol.

A.19 Nagios Service

nagios.xsede.org at IU. Contact: Joe Rinkovsky <jrinkovs@iu.edu>. Nagios is a user-level, centralized monitoring system that runs periodic tests to determine the state of XSEDE provided service in particular, services that level 1 & 2 SPs have registered with XSEDE as being provided. This data is used to track uptime and by the XOC to contact service owners in order to resolve outages.

Data

There is no sensitive data on these system, just statistics like uptimes.

There are some stored credentials on the Nagios service for unprivileged accounts.

Data Flow

Authentication to the Nagios server is via Kerberos.

All web traffic is over HTTPS. Some services have credentials saved in their configuration files but these credentials are only sent over the network via encrypted channels.

System Interfaces

Administrators login through the web interface over HTTPS.

System configuration is managed via SSH. Logins are via two-factor authentication.

Nagios polls systems over HTTPS.

A.20 OAuth

oa4mp.xsede.org at IU. Contact: Ed Berger <eberger@psc.edu>. The XSEDE OAuth Service authenticates XSEDE users and issues short-lived certificates to XSEDE science gateways (web portals), such as Globus Online. For more details see https://portal.xsede.org/oauth/.

Data

The XSEDE OAuth Service maintains a database of registered gateways/portals and their associated public keys. No passwords are stored.

Data Flow

Users enter their XSEDE username/password on a web form when authenticating to this service.

System Interfaces

All interactions are over HTTPS (port 443) except for the MyProxy communications over TLS.

A.21 Secure Jabber Server

At NCSA. Contact: Warren Raquel <wraquel@ncsa.illinois.edu>. This is an encrypted chat service used by the XSEDE Incident Response Team.

Data

Discussions of XSEDE security incidents are facilitated through this server, but it does not archive conversations.

Data Flow

Chat messages flow between chat clients on the IRT member systems and the chat server.

System Interfaces

Jabber (XMPP) protocol over TLS. TLS is forced.

A.22 perfSONAR

HOSTNAME	LOCATION	ADMIN	COMMENTS
ps.<sitename>.xsede.org	Each level 1 SP	Kathy Benninger <benninger@psc.edu>	
psarch.psc.xsede.org	PSC		

perfSONAR provides an active, distributed network test infrastructure for XSEDE with one pS Measurement Point (MP) deployed at each level 1 SP and one central Measurement Archive (MA) display server deployed at PSC. The pS MPs run a mesh of scheduled BWCTL iperf, OWAMP, and traceroute tests between SP sites with results collected by an esmond monitoring daemon and offered via web interface. The MPs are available to members of the XSEDE ops-network group as network test points for on-demand testing (e.g. iperf, owping, tcpdump) when more in-depth analysis is required. Each MP also offers, reverse ping, reverse traceroute, NPAD and NDT services via a web interface along with NPAD and NDT access via command line.

Data

The automatic data collected are network statistics from OWAMP, traceroutes, and IP performance testing tools (iperf3). None of these archives are sensitive.

On-demand reverse pings and reverse traceroutes can be performed at any MP. Also, tcpdump traces can be collected at any one point (by authorized users) which may be more sensitive. The output of these on-demand services is not regularly archived to the MA.

Data Flow

Measurement data flows from each of the MPs to the MA. This data is not sensitive and is not encrypted in transit.

System Interfaces

Management of the system can be done over HTTPS or on the command line through SSH. Reverse ping, reverse traceroutes, NDT, and NPAD are available over HTTP to anyone.

The MA will run a collection service to store data from all sites in a sql database.

A.23 Resource Description Repository

rdr.xsede.org at PSC. Contact: Ed Hanna <ehanna@psc.edu>. Resource Description Repository is a common, central repository where Service Providers (SP) will publish new, and manage existing, resource descriptions. RDR provides the interface for SPs, and other entities, to publish and manage resource descriptions.

Data

The RDR data is stored in the the teragrid database in the rdr schema in the production server as the XSEDE Central Database (XDCDB) at SDSC. It contains compute, grid, storage and data collection resources along with conversion factors. A complete history of all resources is also maintained.

Data is not sensitive, but integrity is important as it contains charge rates and conversion factors.

Data Flow

SPs update the web interface at rdr.xsede.org with new and existing resource information. rdr.xsede.org provides an xml interface to the data that allows the data for an SP's resources to be downloaded by IIS kits at each SP. From the local IIS client the data is published into IIS and the portal retrieves any needed data through IIS interfaces.

System Interfaces

The end result of the RDR data path is available from https://www.xsede.org/resources/overview. On the backend data is either moved over web services (HTTPS) or a JDBC over SSL connection to the postgresql server at SDSC.

A.24 RSA SecurID

rizzo.ncsa.utk at NICS. Contact: Gary Rogers <grogers3@utk.edu>. The XSEDE RSA SecurID service is to allow admins to use an XSEDE domain RSA server for administering central services. Some use local two-factor services in addition or instead of the XSEDE one.

Data

The RSA servers contain usernames and hashed passwords for all XSEDE users. It also contains cryptographic seeds for the tokens.

Data Flow

RSA servers are either directly contacted by services for authentication or through RADIUS servers.

System Interfaces

Administration is done through a web browser over HTTPS. RSA ships a virtual appliance that is mostly a black box. Other interaction happen over proprietary protocols or RADIUS.

A.25 Ticket System

tickets.xsede.org at TACC. Contact: Mike Packard <mpackard@tacc.utexas.edu>. The XSEDE Ticket System provides XSEDE an interface to report problems and a standard, XSEDE-wide method of tracking and resolving those problems. The ticket contains the complete record of responses, replies, and notations generated by the resolution of the problem. XSEDE uses RT (Request Tracker) for its ticket system interface. There is a failover backup system at NICS that can be switched to primary within a few moments, and without any data loss.

Data

The ticket system data is stored in a PostgreSQL relational database system at TACC This data includes all of the email and log entry text fields along with the metadata for each ticket. Sometimes this may include sensitive information such as reports of security incidents. However, there is little in regards to PII except the full name of users.

Data Flow

Email sent to specific XTS addresses are routed directly to the ticket system, which automatically creates a ticket. User responses are automatically added it to an existing ticket.

Newly created tickets are assigned to the XSEDE Operations Center. From there the ticket is evaluated and then assigned to a group/individual with whom the problem will be resolved. Actions on tickets may also generate emails back to the original reporter.

The other data flow is between the web front end and the backend database system of which this is one database.

System Interfaces

The ticket system is accessible directly via https://tickets.xsede.org/ or through the User Portal at https://portal.xsede.org/group/xup/tickets. The other front-end interface is through receipt of email to help@xsede.org, which generates a new ticket.

A.26 Sciforma

projects.xsede.org at IU. Contact: Scott Simmerman <simmerma@eecs.utk.edu>. Sciforma is the project scheduling tool that is used to create the schedules for quarterly and annual reports as well as the planning schedule. In addition the ECSS areas use is to manage the ECSS resources and projects. Sciforma is available to all areas of XSEDE that need a project scheduling capability. Some areas also use the staff and staff allocations information to populate their area wiki page staff directory.

Data

The data contained in Sciforma is basic XSEDE staff information much of which is also available on the XSEDE wiki or on their institution's staff directory. Their is nothing confidential about the project information and many times we end up providing images of the project schedules on the wiki or emailing them to the XSEDE staff.

Data Flow

The only automatic data flow out of Sciforma is to populate the wiki pages for some WBS area's staff directories. All of the information input into Sciforma is all manually entered. The current version of Sciforma uses JAVA and a PostgreSQL DB. Sciforma has a XSEDE built authentication component that allows it to be secured using the XUP Kerberos ID/PW.

System Interfaces

Sciforma has a client interface which has the most capabilities and is what everyone is told about and is directed to use. Sciforma also has a web interface version that has more limited capabilities and we do not generally recommend that it is used. For manual backups and restores of the DB Sciforma provides an admin web page. Sciforma is backed up automatically on a dailty basis.

A.27 Security Wiki

https://ops-security.xsede.org/ at NCSA. Contact: Warren Raquel <wraquel@illinois.edu>. This is a collaborative workspace for incident response.

Data

The data is very sensitive having to do with current and past security events for XSEDE. Access is limited to XSEDE security and incident response team members.

Data Flow

Data does not flow except through the file attachment interface of the wiki.

System Interfaces

Media wiki server running on HTTPS.

A.28 SharePoint

NAME	WHERE	EMAIL	SHAREPOINT RESPONSIBILITIES
Michael Gates	SDSC	mgates@sdsc.edu	Administrator
Michael Northrop	UC	mikenorthrop@uchicago.edu	Security Administrator

A Microsoft wiki used for quarterly and annual reporting, planning for future years, and for some special XSEDE groups to collaborate on documents.

Data

This website is reserved for groups of users collaborating on XSEDE reporting. It's private because only certain people are supposed to edit it, everything entered eventually becomes a part of public reports. So nothing is very confidential.

Data Flow

It is a self-contained system at SDSC. However, XSEDE is not the only project with a virtual space within this share point system. Data does not flow between spaces though. While it is possible to create connections to the Sharepoint data without using the web interface, these methods must be setup and configured properly to work. There are currently no such alternate access methods existing to our Sharepoint environment. Sharepoint uses an MSSQL database for its backend, it is possible to access the database via SQL tools in the way that you would any database. However, our SQL server is limited to specific services each with their own service account and access only to their own databases and the IT systems administrators at the site.

System Interfaces

All interaction is over HTTPS, authentication backed by active directory. Service account connections access only their databases with unique creditials.

A.29 Single SignOn Hub

login.xsede.org at IU. Contact: David Carver <dcarver@tacc.utexas.edu>. The XSEDE SSO (Single Sign-on) Hub SSH service logs XSEDE users in after verifying their XSEDE kerberos/portal username/password with the XSEDE MyProxy service. On successful authentication, the MyProxy service issues short-lived certificates which are made available to the user in /tmp on the hub. Once logged into the hub, a user will be able to gsissh into applicable XSEDE GSISSH service nodes using the short-lived certificate with GSI authentication. For more details
see https://portal.xsede.org/web/xup/single-sign-on-hub.

Data

The XSEDE SSO Hub service doesn't maintain any passwords for portal group users. It does receive (at login time) and store in /tmp, short-lived certificates for users that have logged in. The short-lived certificates are protected read-only for the specific users they belong to.

Data Flow

Users enter their XSEDE username/password when authenticating to this service. The service in turn communicates the users' XSEDE username/password to the XSEDE MyProxy service. On successful authentication, the XSEDE MyProxy service issues short-lived certificates for the users' use on the hub.

System Interfaces

All interactions are over SSH (port 22) except for the MyProxy communications over TLS and gsissh communications over custom ports used by certain XSEDE GSISSH services.

A.30 Software Distribution & Source Repository

https://software.xsede.org/ at ANL. Contact: Lukasz Lacinski <lukasz@ci.uchicago.edu>. The SVN-based source repository and software packages for production, campus bridging, gateways, security, and development are hosted on this web server run at the IU web farm. Use cases used by the architecture team and a deprecated XSEDE Bugzilla service are also hosted on this web server. While it has the look and feel of the XUP and website on the main page, it is a separate server managed by separate teams.

Data

Data includes software for SPs, the certificate trust store and many documents under revision control. The integrity of the software distributed is of course vital. None of the data is sensitive.

Data Flow

Users can checkout software to their client systems with SVN or just browse through the web interface. Only a few dozen people are able to write to these repositories. Authentication is via XSEDE Kerberos.

System Interfaces

SVN over SSH or HTTP/S.

A.31 Speedpage

speedpage.psc.edu at IU. Contact: Bob Budden <rbudden@psc.edu>. Speedpage is an analysis tool used to test the performance of data movement to/from XSEDE resources. It uses the GridFTP protocol to initiate and monitor 3rd party copies between sites. Data is kept in a MySQL server and published via an Apache web interface.

Data

All of the data obtained by the Speedpage is world viewable via the web site or MySQL read only access. The system itself is protected by UNIX file permissions. However, a Kerberos keytab is kept for automation of the service. Short term, 24 hr, MyProxy certificates for and are obtained via the keytab to use for grid authentication.

Data Flow

Data is collected by issuing 3rd party globus-url-copy commands and parsing the output. Data is then stored in a MySQL database and published via Apache web server.

System Interfaces

SSH	Limited number of PSC only staff access
MySQL	Read only account for XSEDE use
Apache	Web interface, world viewable

A.32 User Profile Service

https://info.xsede.org:8444/web-apps/html/profile-v1/usage at IU. Contact: Rion Dooley <dooley@tacc.utexas.edu>. The User Profile Service is an authenticated set of endpoints that provide user-specific views of XSEDE systems and services. While the root domain is collocated with the IIS services, it runs as a separate service in its own Tomcat container on each of the IIS services. The User Profile Service relies heavily on the XSEDE central database, a CouchDB instance, GPIR, and direct interaction with remote systems for its information. While every effort has been made to make the service interactions similar to the IIS services, the User Profile Service is a separate project developed and maintained outside of IIS.

Data

Data is primarily gathered at invocation time from the XSEDE central database and CouchDB. Some data is cached on disk. This data is protected by standard UNIX permissions. None of the data collected is highly sensitive though can reflect the activity of user jobs.

Data Flow

Consumers pass their XSEDE username/password to the service using HTTPS Basic authentication on each invocation. The resulting data is passed via HTTPS back to the consumer in one of several user-specified formats.

System Interfaces

All forward facing interfaces are accessed over HTTPS (port 4443). Access to the XSEDE central database is done via SSL. Communication with the remote MyProxy service is over TLS. Communication with the CouchDB instance is via HTTPS.

A.33 Central Database

HOSTNAME	LOCATION	ADMIN	COMMENTS
balthazar.sdsc.edu	SDSC	lcarson@sdsc.edu	
franco.psc.edu	PSC	Rob Light <light@psc.edu>	

tgcdb.xsede.org points to one of the above two hosts.

The primary instance of the XDCDB runs at SDSC and the failover instance runs at PSC. At SDSC, the following servers are run: (**what is run at psc**)

DESCRIPTION	HOST	PORT	DATABASE
production	tgcdb.xsede.org	5432	teragrid
testing	tgcdb.xsede.org	3333	teragrid
development	tgcdb.xsede.org	3333	tgcdb_test

The XDCDB is a PostgreSQL relational database that contains information regarding all XSEDE projects, allocations, users, resources and usage.

The XDCDB consists of a number of schemas. The primary schema is the "acct" schema, which contains all of the project-related data. Other important schemas include the "portal" schema, used by the XSEDE User Portal (XUP) and the "amie" schema, used by the Account Management Information Exchange (AMIE) data transport system.

Data

A complete list of XDCDB schemas is below:

SCHEMA	DESCRIPTION	LOCATION
acct	User Account Database	
amie	Account Management Information Exchange	http://svc.bu.edu/AMIE
gateway	XSEDE Science Gateway	https://www.xsede.org/web/guest/gateways-listing
gxmap	Globus	https://www.xsede.org/globus-online
info_services		
mcp	Master Control Program	https://www.xsede.org/metascheduling

myproxy	GSI=SSH tools	https://www.xsede.org/oauth
xras	XSEDE Resource Allocation System Review Application	https://xras-review.xsede.org and https://xras-admin.xsede.org
portal	XSEDE User Portal	https://portal.xsede.org
pra	Person Reconciliation Application	internal
rdr	Resource Description Repository	https://portal.xsede.org/web/guest/resources/overview
user_services	XSEDE User Services	https://www.xsede.org/user-services

PII Notes: The XDCDB stores the full name of users as given to us.

Data Flow

The primary means of data flow both into and out of the XDCDB is via **AMIE.** Data for projects, allocations, and PI/co-PI accounts originates in the XSEDE Resource Allocation System (XRAS), and is transported via AMIE to the XDCDB. Other account requests originate in the XUP, inserted directly into the XDCDB. The XDCDB in turn sends the data via AMIE to the relevant SPs. SPs send usage back to the XDCDB also via AMIE.

System Interfaces

The system interfaces to the XDCDB are XRAS, AMIE, the XUP and the xdusage command.

A.34 XRAS Service

HOSTNAME	LOCATION	ADMIN	COMMENTS
xras-review.xsede.org	NCSA	aschuele@illinois.edu	
xras-admin.xsede.org	NCSA	aschuele@illinois.edu	

XRAS (XSEDE Resource Allocation System) provides an interface for request review and administration for XSEDE resources. The system consists of two Web interfaces for review and administration. The backend of these Web interfaces is the XSEDE Central Database (XDCDB). There are also several supporting services: an identity service, an accounting service, and a rules engine as well as a RESTful API to communicate with the submission interface located in the XSEDE User Portal (XUP).

Data

XRAS data is stored in the XDCDB. This data includes all of the information contained in the individual request submitted to XRAS, along with the data associated with resources, opportunities, reviews and reviewers.

PII Notes: XRAS stores the full name of users as given to us.

Data Flow

The request submission data is entered into the XUP. Data is stored in the XDCDB via the XRAS API.

Administrators have access to all data in the request process through an administration interface. This interface allows administrators to modify requests, assign reviewers, award requests, and initiate the data flow process to the XSEDE central database (XDCDB). This transfer of data is via the AMIE.

Administrators also prepare reports for the periodic allocations review board meetings via the administration interface.

Reviewers periodically access XRAS via the reviewer interface. Reviewers are allowed to view the current and past requests and reviews to which they are assigned along with past usage of associated requests. They then submit a review in text format of the request along with a numeric suggestion.

System Interfaces

XRAS is accessible through the XSEDE User Portal for Proposal Submission https://portal.xsede.org/submit-request .

Reviewers access XRAS via https://xras-review.xsede.org/ Administrators access XRAS via https://xras-admin.xsede.org/ .

The supporting services are located at: Identity Service: https://xras-identity-service.xsede.org/v# Accounting Service: https://xras-accounting-service.xsede.org/v# Rules Engine: https://xras-rules-service.xsede.org/api/v#/teragrid API: https://xras-submit-api.xsede.org/v# *where v# indicates the current release version of the supporting service*

XRAS communicates with the XDCDB via AMIE, the accounting service and the identity service.

A.35 User Portal, KnowledgeBase, Web Site, and Wiki

HOSTNAME	LOCATION	ADMIN	COMMENTS
www.xsede.org	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
xdsecure.xsede.org	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
mobile.xsede.org	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
portal.xsede.org	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
xupdb.tacc.utexas.edu	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
xsede-httpd.tacc.utexas.edu	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
xuplogin.tacc.utexas.edu	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	
api.xsede.org	TACC	Maytal Dahan <maytal@tacc.utexas.edu>	

The webpage, user portal and staff wiki are all presented as one unified environment to users, though the different pieces run on different software platforms with different access control lists. Furthermore, different functions such as login and the database are on different hosts. While not 100% critical to using XSEDE, the portal is the public face of the project and thus very important.

Data

Most of the private content is on the staff wiki, which is sensitive and password protected. XUP users will also need to login to manage their profile and certificate DNs. Certificates and other credentials are no longer stored on the server.

Data Flow

Some requests will push data to the XRAS system and from there to the XDCDB through AMIE. Data is also going more directly back and forth from the XDCDB through the AMIE server.

There is also Kerberos communications with the XSEDE KDC and the xdsecure system.

Finally, the XUP communicates with the IIS and knowledgebase over web services (HTTP/REST).

System Interfaces

There is of course the traditional HTTP and HTTPS interface used by most users and between the different servers. Mysql over SSL is used to connect to the backend database. Also, on the backend there is communication over the AMIE and the Kerberos protocol. SSH through two factor is used for OS system administration.

Appendix B. Risk Assessment Rankings

Initial Risk Rankings Votes Matrix

	Basney	Butler	Hazlewood	Rogers	Schule	Simmel	Slagell	Sparks	Average	Rounded	STD
Vulnerability Index											
1	1	1	0	1		2	0	0	0.71	1	0.69985421
2	1	1	1	2		0	1	0	0.86	1	0.63887656
3	0	1	1	1		0	0	0	0.43	0	0.49487166
4	0	1	1	2		1	0	0	0.71	1	0.69985421
5	0	0	1	2		0	1	1	0.71	1	0.69985421
6	**0**	**0**	**0**	**2**		**0**	**1**	**0**	**0.43**	**0**	**0.72843136**
7	1	0	1	1		1	0	0	0.57	1	0.49487166
8	1	0	0	1		0	0	1	0.43	0	0.49487166
9	1	1	1	1		2	0	1	1.00	1	0.53452248
10	1	1	1	2		1	2	1	1.29	1	0.45175395
11	0	0	0	1		1	1	0	0.43	0	0.49487166
12	0	1	1	2		1	0	0	0.71	1	0.69985421
13	2	1	1	2		1	1	1	1.29	1	0.45175395
14	0	1	0	1		0	1	2	0.71	1	0.69985421
15	**1**	**0**	**0**	**2**		**2**	**0**	**2**	**1.00**	**1**	**0.9258201**
16	**1**	**1**	**2**	**2**		**2**	**0**	**2**	**1.43**	**1**	**0.72843136**
17	**0**	**0**	**0**	**2**		**1**	**0**	**2**	**0.71**	**1**	**0.88063057**
18	1	1	1	2		1	1	2	1.29	1	0.45175395
19	0	1	0	1		0	0	1	0.43	0	0.49487166
20	0	1	0	1		1	0	1	0.57	1	0.49487166
21	0	0	1	1		0	0	0	0.29	0	0.45175395
22	1	1	1	2		1	1	2	1.29	1	0.45175395
23	0	1	1	2		0	1	0	0.71	1	0.69985421
24	**0**	**1**	**0**	**2**		**0**	**0**	**0**	**0.43**	**0**	**0.72843136**
25	1	1	1	2		2	1	1	1.29	1	0.45175395
26	**0**	**1**	**2**	**2**		**2**	**0**	**1**	**1.14**	**1**	**0.83299313**
27	1	2	2	2		1	2	2	1.71	2	0.45175395
28	1	2	0	2		1	1	1	1.14	1	0.63887656
29	1	2	2	1		2	2	2	1.71	2	0.45175395
30	0	1	0	1		1	1	2	0.86	1	0.63887656
31	1	1	2	1		2	2	2	1.57	2	0.49487166
32	1	2	0	1		0	1	0	0.71	1	0.69985421
33	1	1	1	2		1	1	2	1.29	1	0.45175395
34	**1**	**2**	**0**	**1**		**2**	**1**	**0**	**1.00**	**1**	**0.75592895**
35	**1**	**2**	**0**	**1**		**2**	**2**	**0**	**1.14**	**1**	**0.83299313**
Enter 0 for low, 1 for medium and 2 for high.											
Reference the table at https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Vulnerability+Identification											
If you are uncomfortable ranking an item, you may leave it blank.											

Final Votes with Phone Meeting Notes

	Basney	Butler	Hazlewood	Rogers	Schule	Simmel	Slagell	Sparks	Average	Rounded	STD
Vulnerability Index											
1	1	1	0	1		2	0	0	0.71	1	0.69985421
2	1	1	1	2		0	1	0	0.86	1	0.63887656
3	0	1	1	1		0	0	0	0.43	0	0.49487166
4	0	1	1	2		1	0	0	0.71	1	0.69985421
5	0	0	1	2		0	1	1	0.71	1	0.69985421
6	**0**	**0**	**0**	**2**		**0**	**1**	**0**	**0.43**	**0**	**0.72843136**
7	1	0	1	1		1	0	0	0.57	1	0.49487166
8	1	0	0	1		0	0	1	0.43	0	0.49487166
9	1	1	1	1		2	0	1	1.00	1	0.53452248
10	1	1	1	2		1	2	1	1.29	1	0.45175395
11	0	0	0	1		1	1	0	0.43	0	0.49487166
12	0	1	1	2		1	0	0	0.71	1	0.69985421
13	2	1	1	2		1	1	1	1.29	1	0.45175395
14	0	1	0	1		0	1	2	0.71	1	0.69985421
15	**1**	**0**	**0**	**2**		**2**	**0**	**2**	**1.00**	**1**	**0.9258201**
16	**1**	**1**	**2**	**2**		**2**	**1**	**2**	**1.57**	**2**	**0.49487166**
17	**0**	**0**	**0**	**2**		**1**	**0**	**2**	**0.71**	**1**	**0.88063057**
18	1	1	1	2		1	1	2	1.29	1	0.45175395
19	0	1	0	1		0	0	1	0.43	0	0.49487166
20	0	1	0	1		1	0	1	0.57	1	0.49487166
21	0	0	1	1		0	0	0	0.29	0	0.45175395
22	1	1	1	2		1	1	2	1.29	1	0.45175395
23	0	1	1	2		0	1	0	0.71	1	0.69985421
24	**0**	**1**	**0**	**2**		**0**	**0**	**0**	**0.43**	**0**	**0.72843136**
25	1	1	1	2		2	1	1	1.29	1	0.45175395
26	**0**	**1**	**2**	**2**		**2**	**0**	**1**	**1.14**	**1**	**0.83299313**
27	1	2	2	2		1	2	2	1.71	2	0.45175395
28	1	2	0	2		1	1	1	1.14	1	0.63887656
29	1	2	2	1		2	2	2	1.71	2	0.45175395
30	0	1	0	1		1	1	2	0.86	1	0.63887656
31	1	1	2	1		2	2	2	1.57	2	0.49487166
32	1	2	0	1		0	1	0	0.71	1	0.69985421
33	1	1	1	2		1	1	2	1.29	1	0.45175395
34	**1**	**2**	**0**	**1**		**2**	**1**	**0**	**1.00**	**1**	**0.75592895**
35	**1**	**2**	**0**	**1**		**2**	**2**	**0**	**1.14**	**1**	**0.83299313**
Enter 0 for low, 1 for medium and 2 for high.											
Reference the table at https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Vulnerability+Identification											
If you are uncomfortable ranking an item, you may leave it blank.											
Changes after call:											
Adam Changed his vote on #16 which removed the outlier.											
Agreement on most items to go with the average. There was a question whether or not to generalize the GO item.											
Call attendees were: Randy Butler, Jim Basney, Gary Rogers, Amy Schule, Adam Slagell, Benny Sparks, & Jim Marstellar											

Final Risk Assessment Votes

	Basney	Butler	Hazlewood	Rogers	Schule	Simmel	Slagell	Sparks	Average	Rounded	STD	
Vulnerability Index												
1	0	0		0	0	1	0	1	0.29	0	0.45175395	
2	0	1		1	0	0	0	1	0.43	0	0.49487166	
3	0	0		0	0	0	0	1	0.14	0	0.34992711	
4	0	1		0	0	1	0	1	0.43	0	0.49487166	
5	0	1		0	0	0	1	1	0.43	0	0.49487166	
6	0	0		0	0	0	1	0	0.14	0	0.34992711	
7	0	0		0	0	1	0	1	0.29	0	0.45175395	
8	0	0		1	0	0	1	0	0.29	0	0.45175395	
9	1	1		0	0	1	1	1	0.71	1	0.45175395	
10	0	0		0	0	1	0	1	0.29	0	0.45175395	
11	0	0		0	0	1	1	0	0.29	0	0.45175395	
12	0	2		1	0	1	0	1	0.71	1	0.69985421	
13	0	1		0	0	1	1	1	0.57	1	0.49487166	
14	0	1		1	0	0	1	1	0.57	1	0.49487166	
15	0	2		0	1	1	1	0	0.71	1	0.69985421	
16	**0**	2		**0**	**0**	**1**	**2**	**2**	**1.00**	**1**	**0.9258201**	**No one obje**
17	0	1		1	0	0	1	1	0.57	1	0.49487166	
18	1	1		0	0	1	1	1	0.71	1	0.45175395	
19	0	0		0	0	0	0	0	0.00	0	0	
20	0	0		0	0	1	0	1	0.29	0	0.45175395	
21	0	1		0	0	0	1	1	0.43	0	0.49487166	
22	1	1		0	0	1	1	1	0.71	1	0.45175395	
23	1	1		0	1	0	1	1	0.71	1	0.45175395	
24	1	1		0	0	0	1	0	0.43	0	0.49487166	
25	1	1		0	0	1	1	0	0.57	1	0.49487166	
26	**Retired risk. Please Skip**								**#DIV/0!**	**#DIV/0!**	**#DIV/0!**	
27	1	1		1	0	1	1	1	0.86	1	0.34992711	
28	1	1		1	0	1	2	2	1.14	1	0.63887656	
29	0	1		0	0	1	1	2	0.71	1	0.69985421	
30	**0**	1		**0**	**0**	**0**	**1**	**2**	**0.57**	**1**	**0.72843136**	**No one obje**
31	1	1		0	1	1	1	2	1.00	1	0.53452248	
32	0	1		0	0	1	1	1	0.57	1	0.49487166	
33	1	1		0	0	1	1	1	0.71	1	0.45175395	
34	0	0		0	1	1	1	1	0.57	1	0.49487166	
35	0	0		0	0	1	1	1	0.43	0	0.49487166	
Enter 0 for low, 1 for medium and 2 for high.												
Index maps to the table at https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Vulnerability+Identification												
If you are uncomfortable ranking an item, you may leave it blank.												
Definitions of high, medium & low are at https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE%20Federation%20Risk%20Assessment#section-XSEDE+Federation+Risk+Assess												
Note: It is hard to argue that something affecting only a few individual accounts can have anything but a low impact by these defnitions.												

Appendix C: XSEDE Relevant Threat Matrix

THREAT-SOURCE	MOTIVATION	THREAT ACTIONS
Cracker	Challenge Ego Rebellion *Activism*	**Credential Harvesting** Web Site Defacement **Privilege Escalation** Malware/**Trojan Insertion** Social Engineering *Cracking Passwords on HPC* Proxy to Hide Behind **Dictionary Attack Services** Establish Foothold / Monitor Trusted XDNet
Computer Criminal	**Monetary Gain** High-powered Attack Platform	**Bitcoin Mining** Launch DoS Social Engineering **Credential Harvesting** **Privilege Escalation** Malware/**Trojan Insertion** Proxy to Hide Behind **Bot-herding** **Criminal Activities (spamming, host illegal content, etc)** **Dictionary Attack Services**
Insiders (Employees or Users)	**Monetary Gain** **Curiosity** Disgruntled **Unintentional Error** (Negligent, poorly trained, programming or configuration error)	**Bitcoin Mining** **Browsing Others' Data** Blackmail *System Sabotage* *Data Corruption/Deletion* **Privilege Escalation** Malware/**Trojan Insertion** **System Bugs** Interception/Monitoring

Bold for common OR expected.

Italic for not seen AND unexpected.

Appendix D: Threat Profile Survey

SCOPE: Keep in mind when answering each question that we are just interested in threats that can affect XSEDE as a whole by breaking the fundamental trust fabric, interrupting shared services, or bringing multiple SPs offline. We are not looking at the level of site-specific threats such as environmental or physical security problems.

An example threat source would be a cyber criminal motivated by monetary gain. A corresponding threat action could be hijacking HPC(s) for bitcoin mining.

#1 What kind of threat sources (e.g., criminals, crackers, hacktivists, insiders) have you seen in Teragrid incidents? What is most common?

NCSA: Mostly hackers for fun. We likely have had criminals that want to use the resources for $$, but most of the incidents that we have caught have been within less than 24 hours, so there has been no time to monetize the hacks.

NICS: It's difficult to determine identity behind an attack, especially if it is caught and thwarted quickly. Hackers that know what they are doing either proxy their connections through an intermediary to hide their identity or use some else's credentials to perpetrate the incident. By far, the most common threat has been account compromise and dealing with "unknown" users on the machines masquerading as legitimate users. Often times the attack sources tend to be from outside the United States. I doubt we have any/many "insider" threats (i.e., TG/XSEDE staff hacking our own systems).

PSC: Generally, we don't know the motivations of attackers responsible for system or account compromises. Mostly we find their main interest is trying to expand the compromise by replacing SSHD, copying Keys and/or adding keys. To my knowledge there have been no known cases of insider hacking or actions by hackivists against TG/XSEDE resources. The most common security related activity is compromised user accounts. I would say about 85% of our effort is dealing with compromised accounts. The source of compromise is usually at the user's home institution, not on an XSEDE system.

Purdue: I'd say criminals and crackers are the most common, with criminals the more common of the two.

SDSC: Of the provided list, we've seen criminals, crackers, and insiders. Aside from a single incident (definitely a cracker), it's difficult to distinguish between criminals and crackers, since we don't know what their motives were. The insider incident(s) appear to be unintentional, rather than out of malice. By far, however, the most common threat has been automated attacks -- massive amounts of password guessing, port scanning, and the like.

#2 Are there any new threat sources you would expect with XSEDE that you have not seen, and if so, why?

NCSA: Not that I can think of.

NICS: Any new threat sources would probably be due to globally changing threat profiles, not anything XSEDE specific. I predict that hactivism will increase, but this is not XSEDE specific.

PSC: One of the goals of XSEDE is to reach more users. More users will likely lead to more compromised accounts.

Purdue: I would expect hacktivism to become a threat due to the increased notoriety that groups like Anonymous and LulzSec and have gained in the past year and the increasing political unrest globally.

SDSC: With the recent media attention on groups like Anonymous and Lulzsec, I expect to see a rise in hacktivism. I don't see XSEDE resources being targeted by these specific groups, as the sort of work we do does not appear to be on their political agenda. However, this type of hacktivism may inspire other groups. I also expect to see a slightly different type of threat, created by outsiders, but launched by unwitting, non-specifically targeted insiders (e.g. trojaned or vulnerable software). The barriers to entry for programming are coming down and more junior and amateur programmers will be creating code that does something our users want to do. These types of programmers are likely to not have the discipline to write good code or protect their distribution chain against 3rd party modifications.

#3 In instances in which you could infer, what has typically motivated (consider negligence a kind of motivation) attacks in Teragrid?

NCSA: Most have been hacks of opportunity or convenience. Credentials were stolen that eventually led to our systems. I can't say that any of those attacks have been targeted specifically for NCSA or the Teragrid.

NICS: The most common motivations are likely: Trophy/bragging; Ease of attack (due to unpatched systems, etc.). Although I can't point to any specific instances, I'm sure compromised TeraGrid systems have been used to send spam and attack other institutions.

PSC: If you're meaning what was the situation that allowed an attack, poorly maintained systems (not properly patched against vulnerabilities), shared passwords that were compromised elsewhere, weak system passwords, improperly configured systems. Note that most of this applies to non-XSEDE systems though we have had some XSEDE compromises due to system vulnerabilities.

Purdue: Based on recent events, financial motivation seems to be the most common.

SDSC: As far as successful attacks on our resources go, they have been for ego and/or the result of unintentional error.

#4 Are there any new attacker motivations you would expect with XSEDE that you have not seen, and if so, why?

NCSA: Not that I can think of. Possibly bitcoin, which we have not seen specifically on a TG system yet but could happen.

NICS: Any new attack motivations would probably be due to globally changing threat profiles, not anything XSEDE specific.

PSC: A number of large-scale grids have recently experienced bitcoin mining. TG/XSEDE has not but I would expect we'll see such activity sometime in the future.

Purdue: Based on #2, I'd expect political motivations to become an issue, either by active attack against targets or attempts at wikileaks-style disclosure.

SDSC: None that I can think of at this time. The motivation list appears to cover the things I can think of.

#5 What kind of incidents or threat actions (DoS launch, bitcoin mining, credential harvesting, password cracking, sabotage, interception, etc.) have you seen in TeraGrid? What are the most common?

NCSA: Most common is compromised user account. Outside of that we have only seen a few others: terms of use violation, Web server vulnerability, internal incident (researcher copying data).

NICS: The most common is by far credential harvesting. We also have a fair amount of unauthorized account sharing, which makes security auditing more difficult.

PSC: Known threats: Credential harvesting, system scanning, brute force attacks, spam on websites (very few). Not seen: Bitcoin mining, password cracking (maybe many, many, years ago), sabotage.

Purdue: I'm aware of DoS and bitcoin mining on resources and password cracking or social engineering attacks against user credentials (sometimes at other resources).

SDSC: On our TG systems, we've seen credential harvesting and stone-stepping or using the compromised account and machine to attack other resources. On our non-TG systems, we've also seen unauthorized use of the host to launch DOS attacks, run bots, run irc proxies, run reverse web proxies, host illegal content, and perform network sniffing.

#6 Are there any new threat actions you would expect with XSEDE that you have not seen, and if so, why?

NCSA: Not that I can think of offhand.

NICS: XSEDE plans to interoperate more with campuses, which will lead to more exposure. Our user base may be wider, and their identity may be less known. Additionally, the grid architecture is changing to include more software and technologies that may have inherent vulnerabilities not yet discovered. Most new threat actions would probably be due to globally changing threat profiles, not anything XSEDE specific.

PSC: I think that we'll see threats from the more open nature of the XSEDE program vs Teragrid. An example of this is the short lived, ad-hoc, network connections that are expected (PSU & UMN). Supporting federated authentication will likely increase the number of XSEDE users which may lead to more compromised accounts, so we'll need to develop a response plan with federation partners.

Purdue: I don't foresee any new threat actions, but I would expect a general increase in existing actions.

SDSC: Along with the rise of hacktivism, I would not be surprised to see attacks on our users' research -- impeding research by destroying data, or invalidating (discrediting) research by tampering with it. Identity (and data in general) theft may be a new threat action if our users deliberately or inadvertently use XSEDE resources to process PII, PHI, or politically sensitive data. As a final comment, I think it's only a matter of time before the criminals and crackers who get on to an HPC machine realize what they have access to and start using the system as a legitimate user would, blending into the noise.

Note: no responses were given by NCAR or TACC.

Appendix E: Vulnerability Identification Survey

#1 How does XSEDE monitor for new security patches for software in the baseline software and services document?

Baer: Most of that software comes from xsede partners in the xsede services document. Troy relies on a push from growls like globus,

McNally: Maybe better to ask Troy for stuff in baseline. Stephen's group is more focused on shared systems and services services in the XSEDE Services Master Spreadsheet. For those, they mostly rely on security ops to tell them if there is a new vulnerability.

#2 How do we ensure each SP is up-to-date? Do we help to roll out updates to the SPs?

Baer: We notify SPs, but don't push them. Unless there is a serious security vulnerability and security ops is pushing for it to rollout, SPa update when and if they feel like it. Also, ctss registration is maintained by hand. It is easy to update software and not update your registration. globus will monitors in an automated way and tells Troy about Sps running old versions. CTSS is not terrible convenient, and the information is stale. It would be nice if INCA or something else could go out and grab this info from SPs.

McNally: Ask Troy for more user facing services and deployments at the SPs.

#3 What software in the baseline services document no longer has support or active development?

Baer: There is no clear list of which pieces are not supported anyone.

McNally: This is a better question for Troy.

#4 Are you aware of any security inconsistencies in deployments that could allow an attacker to more easily gain access at one site and use the trust fabric of XSEDE (e.g., GSI certificates) to spread their attack?

Baer: Some sites accept passphraseless ssh keys, some will but don't let users manage them.

Marsteller: More than even inconsistencies in password and authentication policies, the SPs differ greatly in how they protect their borders and harden their systems. Some even allow telnet to through their border, and some don't do any network monitoring. This all ties back to the need for a baseline services document that is agreed upon and followed.

McNally: There is no detailed security baseline. A lot of this is ad hoc. SDSC has quarterly updates of AMIE and XDCDB. More often than not Stephen's group follows the leads of SPs like SDSC.

#5 Would you add any items or classes of items to the exiting vulnerability table?

Baer: Genesis II containers are like their own CAs. Maybe we need to sign genesis root certs. Hopefully new services like this won't require rolling out and maintaining a separate PKI from what we already have.

Marsteller: It is pretty complete without any major categories or areas uncovered.

McNally: POPS has no offsite redundancy, ticket system. Critical pieces seem to all be replicated off site.

Quinn: POPs is not replicated offsite currently. POPS is like Fastlane for XSEDE. It is a separate Sybase database and connects via AMIE to the XDCDB. It is the starting point for all new accounts and allocations.

Wefel: Maybe DoS attacks using XD resources. That is more of a threat than a vulnerability though. Adam can go through the table and see if there is a vulnerability to match that threat against.

#6 Would you modify the vulnerabilities of corresponding threats already listed?

Baer: Need to take a closer look.

Marsteller: Nothing major, but Marsteller will make some minor edits to the wiki page directly.

McNally: None of them are supercritical. Look at the tiers in the master spreadsheet.

Wefel: DNSSEC is used, but not exclusively. We delegate subdomains to other SPs, and for some of those we sync with DNSSEC. The xsede.org root DNS at NCSA and SDSC sync over regular tcp connection to a hidden master server unreachable from other IPs.

#7 Are you aware of security policies or procedures that are not effectively being realized?

Marsteller: Most policies are not being effectively communicated and none audited for compliance. This is especially true for sites new to XSEDE.

#8 What policies are we missing or most critically need to update?

Marsteller: Most policies are out-of-date as they are from TG. The most important to update in the short term is the security baseline document and incident response playbook. We should also create a document explaining the security responsibilities of SPs (at least for level 1 & 2). Another thing to consider is putting in place processes for revoking credentials for key XSEDE staff.

#9 Are you aware of any XSEDE resources or services that have experienced downtime as a result of a security incident?

Marsteller: Not in XSEDE yet, but in TG it was not uncommon for a site's compute resources to go offline because of a compromise. Some other shared services have gone down due to compromise, like INCA, but at least no critical shared services like the portal, CA or central database.

#10 What is the most sensitive kind of data held in the XCDB?

Quinn: There is contact info on people and allocations info for PIs. However, all accepted proposals are public as well as most contact info.

#11 Who can see each type of data?

Quinn: Each SP has an account management person with read-only access. TAS XDmod group (at Buffalo) has read-only access and a synchronized copy of the database. The half dozen DBAs at NCSA & SDSC have both read and write access.

#12 How is the XCDB backed up and how often?

Quinn: It is synchronized between PSC and SDSC every 10 minutes. SDSC backs it up every night.

#13 How do we know a that an XSEDE user in the database maps to a real person and that the contact info is correct?

Quinn: Vetting of PIs is done through the award process by committees, such as, XRAC. Users are vetted when by PIs when they are added to an allocation. Out of band the user gives the PI the account name to add. Users fill out their own contact info and could falsify that, though the PI should at some point know the mapping of person to account name from when they added someone to an allocation. Prior to adding someone to an allocation, there is no record in the XDCDB or DN.

#14 How do we know credentials get to the right person? Can you walk me through the process?

Quinn: Credentials are created with the account creation process. Anyone can create an account, which is nothing more than a Kerberos principal and login to the XUP. They set the password when they create the account.

#15 What is the process for resetting XSEDE passwords?

Quinn: It is self-service on the portal.

#16 What is the process for correcting mistaken information in the XCDB?

Quinn: It is self-service on the portal for user information; manual for things like mistakes in the allocation.

#17 Do accounts become deactivated after a certain amount of time without use?

Quinn: Maytal may have that setup now, but Steve is unsure.

#18 What are the critical services for data services in XSEDE, including supporting services such as CAs?

Simmel: These are mostly the AuthN/Z infrastructure (PKI, Kerberos, local SP accounts) and gridftp services for now. We'll see when GFFS is rolled out. There are local archives like MSS and Golem as well as a wide area Lustre filesystem called Albito. But these are not XSEDE-wide resources, and it's hard to call them critical. Derek really didn't know anything about VFS.

#19 Does XSEDE data services need to be very concerned about PII or proprietary data services?

Simmel: No real PII, in fact you aren't supposed to put it on our resources as an XSEDE user per the AUP. If SPs allow non-XSEDE users special privilege, like private sector partners, that's their own business.

#20 What sort of service guarantees are we giving to users? For example, in terms of data isolation, privacy and backups?

Simmel: We give few if any guarantees on privacy. The default POSIX permissions are sufficient to meet those minimal guarantees. Some sites may promise to backup home directories, but it is not an XSEDE policy

#21 Is the security model consistent regardless of how data sets are moved (e.g., GO vs. XUP files system services)?

Simmel: It depends on what consistent means. But depending on how data is transferred, it may or may not be encrypted. Authentication mechanisms may differ. We are consistent in that all mechanisms meet the minimal guarantees promised users. None of this considers campus bridging.

#22 Does the unavailability of and XSEDENet link fall over to other routes? In other words, would services lose connection if XSEDENet went offline, or just have reduced service?

Wefel: Service would be reduced possibly, but nothing would be unreachable.

#23 Are there privacy concerns if a perfsonar node is compromised?

Wefel: While root on that node can get full packet captures, this is just performance traffic directed to the perfsonar nodes. So one cannot sniff user data or xfers.

#24 Is there any private data in the XSEDE DNS?

Wefel: No, it is just public cnames, A records and SRV records.

#25 Would compromise of the DNS servers allow anything other than DoS attacks, or are there services trusting the DNS responses without other authentication mechanisms?

Wefel: Unsure. Slagell can only think of SSH that is not gsi-based and hence doesn't use certificates. Basney notes "The other DNS insecurity that comes to mind is that Globus software by default relies on DNS for hostname canonicalization.
See:http://dev.globus.org/wiki/C_Security:_Server_Identity_Processing_In_GSI_C. Setting GLOBUS_GSSAPI_NAME_COMPATIBILITY=STRICT_RFC2818 in your environment disables the insecure hostname canonicalization in Globus code."